arXiv:1607.03238v5 [cs.AR] 12 Feb 2017

Scratchpad Sharing in GPUs

Vishwesh Jatala, Indian Institute of Technology, Kanpur
Jayvant Anantpur, Indian Institute of Science, Bangalore
Amey Karkare, Indian Institute of Technology, Kanpur

GPGPU applications exploit on-chip scratchpad memory available in the Graphics Processing Units (GPUs) to improve performance. The amount of thread level parallelism (TLP) present in the GPU is limited by the number of resident threads, which in turn depends on the availability of scratchpad memory in its streaming multiprocessor (SM). Since the scratchpad memory is allocated at thread block granularity, part of the memory may remain unutilized. In this paper, we propose architectural and compiler optimizations to improve the scratchpad memory utilization. Our approach, called *Scratchpad Sharing*, addresses scratchpad under-utilization by launching additional thread blocks in each SM. These thread blocks use unutilized scratchpad memory and also share scratchpad memory with other resident blocks. To improve the performance of scratchpad sharing, we propose *Owner Warp First (OWF)* scheduling that schedules warps from the additional thread blocks effectively. The performance of this approach, however, is limited by the availability of the part of scratchpad memory that is shared among thread blocks.

We propose compiler optimizations to improve the availability of shared scratchpad memory. We describe a scratchpad allocation scheme that helps in allocating scratchpad variables such that shared scratchpad is accessed for short duration. We introduce a new hardware instruction, *relssp*, that when executed, releases the shared scratchpad memory. Finally, we describe an analysis for optimal placement of *relssp* instructions such that shared scratchpad memory is released as early as possible, but only after its last use, along every execution path.

We implemented the hardware changes required for scratchpad sharing approach and the new instruction (*relssp*) using the GPGPU-Sim simulator, and implemented the compiler optimizations in Ocelot framework. We evaluated the effectiveness of our approach on 19 kernels from 3 benchmarks suites: CUDA-SDK, GPGPU-Sim, and Rodinia. The kernels that under-utilize scratchpad memory show an average improvement of 19% and maximum improvement of 92.17% compared to the baseline approach, without affecting the performance of the kernels that do not waste scratchpad memory.

Keywords: Scratchpad Sharing, Thread Level Parallelism, Control Flow Graph, Code Motion

1. INTRODUCTION

The throughput achieved by a GPU (Graphics Processing Unit) depends on the amount of thread-level-parallelism (TLP) it utilizes. Therefore, improving the TLP of GPUs has been the focus of many recent studies [Yang et al. 2012; Kayiran et al. 2013; Anantpur and Govindarajan 2014; Hayes and Zhang 2014]. The TLP present in a GPU is dependent on the number of resident threads. A programmer interested in parallelizing an application in GPU invokes a function, called *kernel*, with a configuration consisting of number of thread blocks and number of threads in each thread block. The maximum number of thread blocks, and hence the number of threads, that can be launched in a Streaming Multiprocessor (SM) depends on the number of available resources in it. If an SM has R resources and each thread block requires R_{tb} resources,

Vishwesh Jatala is supported by TCS Ph.D. fellowship. Jayvant Anantpur acknowledges the funding received from Google India Private Limited.
This article is extension of the paper "Improving GPU Performance Through Resource Sharing", in Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC '16). The paper describes a resource sharing technique that makes architectural modifications to improve GPU performance. This work extends the paper by introducing compiler optimizations to leverage the resource sharing approach.
Author's addresses: Vishwesh Jatala, Department of Computer Science and Engineering, Indian Institute of Technology, Kanpur; Jayvant Anantpur, Supercomputer Education and Research Centre (SERC) , Indian Institute of Science, Bangalore; Amey Karkare, Department of Computer Science and Engineering, Indian Institute of Technology, Kanpur;

then $\lfloor R/R_{tb} \rfloor$ number of thread blocks can be launched in each SM. Thus utilizing $R_{tb} * \lfloor R/R_{tb} \rfloor$ units of resources present in the SM; the remaining R mod R_{tb} resources are wasted. In this paper, we propose an approach, *Scratchpad Sharing*, that launches additional thread blocks in each SM. These thread blocks help in improving the TLP by utilizing the wasted scratchpad memory and by sharing the scratchpad memory with the other resident thread blocks. We further propose *Owner Warp First (OWF)*, a warp scheduling algorithm that improves performance by effectively scheduling warps from the addition thread blocks.

In our experiments we observed that the performance of scratchpad sharing depends on the availability of the scratchpad memory that is shared between the thread blocks. We have developed static analysis that helps in allocating scratchpad variables into shared and unshared scratchpad regions such that the shared scratchpad variables are needed only for a short duration. We modified the GPU architecture to include a new hardware instruction (*relssp*) to release the acquired shared scratchpad memory at run time. When all the threads of a thread block execute the *relssp* instruction, the thread block releases its shared scratchpad memory. We describe an algorithm to help compiler in an optimal placement of the *relssp* instruction in a kernel such that the shared scratchpad can be released as early as possible, without causing any conflicts among shared thread blocks. These optimizations improve the availability of shared scratchpad memory.

The main contributions of this paper are:

(1) We describe an approach to launch more thread blocks by sharing the scratchpad memory. We further describe a warp scheduling algorithm that improves the performance of the GPU applications by effectively using warps from additional thread blocks.
(2) We present a static analysis to layout scratchpad variables in order to minimize the shared scratchpad region. We introduce a hardware instruction, *relssp*, and an algorithm for optimal placement of *relssp* in the user code to release the shared scratchpad region at the earliest.
(3) We used the GPGPU-Sim [Bakhoda et al. 2009] simulator and the Ocelot [Diamos et al. 2010] compiler framework to implement and evaluate our proposed ideas. On several kernels from various benchmark suites, we achieved an average improvement of 19% and a maximum improvement of 92.17% over the baseline approach.

The rest of the paper is organized as follows: Section 2 describes the background required for our approach. Section 3 motivates the need for scratchpad sharing and presents the details of the approach. Owner Warp First scheduling is described in Section 4. Section 5 presents the need for compiler optimizations. The optimizations themselves are discussed in Section 6. Section 7 analyzes the hardware requirements and the complexity of our approach. Section 8 shows the experimental results. Section 9 discusses related work, and Section 10 concludes the paper.

2. BACKGROUND

A typical NVIDIA GPU consists of a set of streaming multiprocessors (SMs). Each SM contains execution units called stream processors. A programmer parallelizes an application on GPU by specifying an execution configuration consisting of the number of thread blocks and the number of threads in each thread block. The number of thread blocks that are actually launched in a SM depends on the resources available in the SM, such as the amount of scratchpad memory, the number of registers. The threads in a SM are grouped into 32 threads, called warps. All the threads in a warp execute the same instruction in SIMD manner. GPU has one or more warp schedulers, which fetch

Table I: Benchmark Applications for which the Number of Thread Blocks is Limited by Scratchpad Memory

	Benchmark	Application	Kernel	#Scratchpad Variables	Scratchpad Size (Bytes)	Block Size
Set-1: Shared scratchpad can be released before the end of the kernel						
1.	RODINIA	backprop	bpnn_layerforward_CUDA	2	9408	256
2.	CUDA-SDK	dct8x8_1 (DCT1)	CUDAkernel2DCT	1	2112	64
3.	CUDA-SDK	dct8x8_2 (DCT2)	CUDAkernel2IDCT	1	2112	64
4.	CUDA-SDK	dct8x8_3 (DCT3)	CUDAkernelShortDCT	1	2176	128
5.	CUDA-SDK	dct8x8_4 (DCT4)	CUDAkernelShortIDCT	1	2176	128
6.	GPGPU-SIM	NQU	solve_nqueen_cuda_kernel	5	10496	64
7.	RODINIA	srad_v2_1 (SRAD1)	srad_cuda_1	6	13824	576
8.	RODINIA	srad_v2_2 (SRAD2)	srad_cuda_2	5	11520	576
Set-2: Shared scratchpad can not be released before the end of the kernel						
9.	CUDA-SDK	FDTD3d	FiniteDifferencesKernel	1	3840	128
10.	RODINIA	heartwall	kernel	8	11872	128
11.	CUDA-SDK	histogram	histogram256Kernel	1	9216	192
12.	CUDA-SDK	marchingCubes (MC1)	generateTriangles	2	9216	32
13.	RODINIA	nw_1	needle_cuda_shared_1	2	8452	32
14.	RODINIA	nw_2	needle_cuda_shared_2	2	8452	32

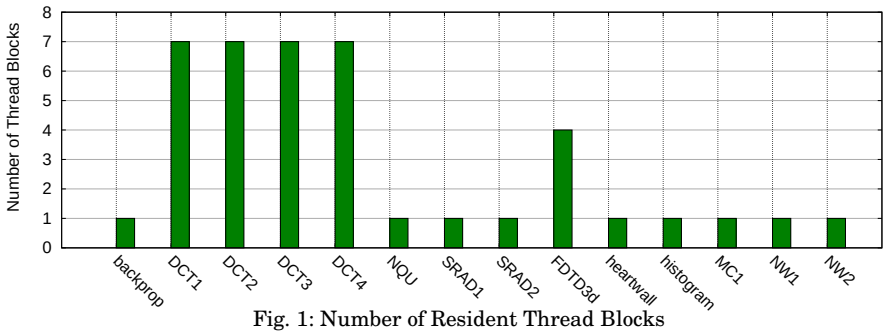

Fig. 1: Number of Resident Thread Blocks

a warp from the pool of available warps based on some warp scheduling algorithm. When no warp can be issued in a cycle, the cycle is said to be a stall cycle.

NVIDIA provides a programming language CUDA [CUDA 2012], which can be used to write an application to be parallelized on GPU. The region of a program which is to be parallelized is specified using a function called kernel. The kernel is invoked with the configuration specifying the number of thread blocks and number of threads as `<<<#ThreadBlocks, #Threads>>>`. A variable can be allocated to global memory by invoking `cudamalloc()` function. Similarly a variable can be allocated to scratchpad memory by specifying `__shared__` keyword inside a kernel function. The latency of accessing a variable from global memory is 400-800 cycles, whereas, latency of accessing from scratchpad memory is 20-30x lower than that of global memory [CUDA 2012].

3. SCRATCHPAD SHARING

Scratchpad memory allocation at thread block level granularity causes scratchpad underutilization. To understand the utilization of scratchpad memory, we analyzed applications shown in Table I. Figure 1 shows the number of thread blocks that are launched in each SM and Figure 2 shows the percentage of unutilized scratchpad memory for the GPU configuration shown in Table II.

Example 3.1. Consider the application *backprop* in Table I. It requires 9408 bytes of scratchpad memory to launch a thread block in the SM. According to the GPU configuration shown in Table II, each SM has 16K bytes of scratchpad memory. Hence

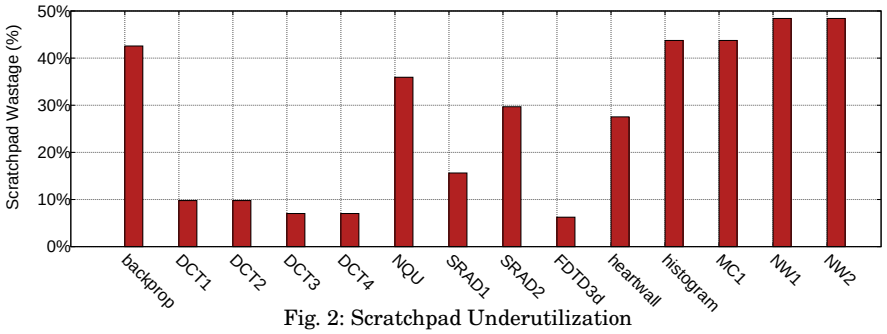

Fig. 2: Scratchpad Underutilization

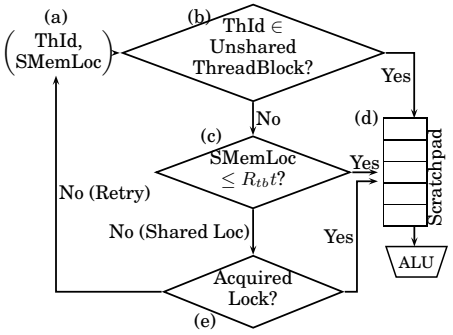

Fig. 3: Scratchpad Access Mechanism

Table II: GPGPU-Sim Architecture

Resource	Configuration
Number of Clusters	14
Number of Cores/Cluster	1
Core Clock	732 MHz
Scratchpad Memory/Core	16KB
Number of Registers/Core	65536
Max Number of TBs/Core	16
Max Number of Threads/Core	3072
Warp Scheduling	LRR
Number of Schedulers/Core	4
L1-Cache/Core	16KB
L2-Cache	1.5MB
DRAM Scheduler	FR-FCFS

only 1 thread block can be launched in the SM, this utilizes 9408 bytes of scratchpad memory. The remaining 6976 bytes of scratchpad memory remains unutilized. We can observe the similar behavior for other applications as well. Hence scratchpad allocation at thread block level granularity not only has lower number of resident thread blocks but also has scratchpad memory underutilization. □

To address this problem, we propose, *Scratchpad Sharing*, that increases the number of resident thread blocks in each SM. These thread blocks use the unutilized scratchpad memory as well as share scratchpad with other resident thread blocks. This not only reduces scratchpad underutilization but also increases TLP on the SM.

Example 3.2. To improve the performance of *backprop*, we launch two thread blocks (say TB_0 and TB_1), which share the scratchpad memory. Instead of allocating 9408 bytes of scratchpad memory to each of TB_0 and TB_1, the scratchpad sharing approach allocates total 16K bytes of memory together for TB_0 and TB_1. In this case, 6976 bytes (the unutilized amount in Example 3.1) of scratchpad memory is allocated to each thread block independently (*unshared* scratchpad), while the remaining 2432 bytes of memory (*shared* scratchpad) is allocated to the thread block which requires it first. For example, if TB_1 accesses the shared scratchpad memory first, it is allocated all of the shared portion. TB_0 can continue its execution till it requires shared scratchpad memory, at which point it waits. TB_0 resumes its execution once TB_1 finishes or releases the shared scratchpad. Thus, TB_0 can help in hiding the long memory latencies of TB_1, thereby improving the run-time of the application. □

To generalize our idea, consider a GPU that has R units of scratchpad memory per SM, and each thread block requires R_{tb} units of scratchpad memory to complete its

execution. Consider a pair TB_0 and TB_1 of shared thread blocks. Instead of allocating R_{tb} units of memory to each of TB_0 and TB_1, we allocate $t \times R_{tb}$ $(0 < t < 1)$ units of scratchpad memory to each of them independently. This is called *unshared* scratchpad. We further allocate $(1-t) \times R_{tb}$ units of scratchpad memory to the pair as *shared* scratchpad. Thus, a total of $(1+t) \times R_{tb}$ units of scratchpad memory is allocated for both. TB_0 and TB_1 can access shared scratchpad memory only after acquiring an exclusive lock, in an FCFS manner, to prevent concurrent accesses. Once a shared thread block (say TB_0) acquires the lock for shared scratchpad memory, it retains the lock till the end of its execution. The other thread block (TB_1) can continue to make progress until it requires to access shared scratchpad memory, at which point it waits until TB_0 releases the shared scratchpad.

The naive scratchpad sharing mechanism, where each thread block shares scratchpad memory with another resident thread block, may not give benefit over default (unshared) approach. we also need to guarantee that in sharing approach, the number of active thread blocks (not waiting for shared scratchpad) is no less than the number of thread blocks in default approach.

Example 3.3. Consider the application DCT3 that requires 2176 bytes of scratchpad memory per thread block. For the given GPU configuration (Table II), 7 thread blocks can be launched in default mode. With scratchpad sharing, it is possible to launch 12 thread blocks (for a certain value of t). Suppose we create 6 pairs of thread blocks where the blocks in each pair share scratchpad. Then, in the worst case, all 12 blocks may request access to the shared portion of scratchpad. This will cause 6 blocks to go in waiting, while only the remaining 6 will make progress. If the shared region is sufficiently large, the application will perform worse with scratchpad sharing.

To make sure at least 7 thread blocks make progress, our approach creates only 5 pairs of thread blocks that share scratchpad memory, the remaining 2 thread blocks are not involved in sharing. Thus, at most 5 blocks can be waiting during execution. □

In our approach, the thread blocks that share the scratchpad memory are referred to as *shared thread blocks*, the rest are referred to as *unshared thread blocks*. The computation of number of shared and unshared thread blocks is described in detail in [Jatala et al. 2016].

To implement our approach, we modify the existing scratchpad access mechanism provided by GPGPUSIM [2014] simulator. Figure 3 shows the scratchpad access mechanism that supports scratchpad sharing. When a thread (Thread Id: $ThId$) needs to access a scratchpad location ($SMemLoc$), we need to check if it is from an unshared thread block. If it belongs to an unshared thread block, it can access the location directly from scratchpad memory (Figure 3 Step (b)). Otherwise, we need to make another check if it accesses unshared scratchpad location (Step (c)). The thread accesses unshared scratchpad location if $SMemLoc < R_{tb}t$ because we allocate $R_{tb}t$ units of scratchpad memory to each of the shared thread blocks. Otherwise, we treat the location as shared scratchpad location. A thread can access unshared scratchpad location directly, however it can access the shared scratchpad location only after acquiring the exclusive lock as shown in Step (e). Otherwise, it retries the access in the next cycle[1].

4. OWNER WARP FIRST (OWF) SCHEDULING

In our approach, each SM contains various types of thread blocks such as, (1) unshared thread blocks, which do not share scratchpad memory with any other thread block, (2) shared thread blocks that own the shared scratchpad memory (*Owner thread blocks*) by having exclusive lock, and (3) shared thread blocks that do not own the shared

[1]The details of required additional storage units are discussed in Section 7.

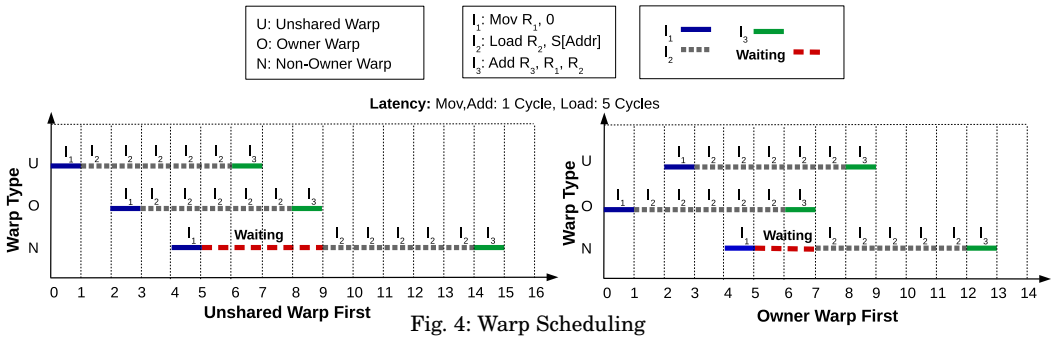

Fig. 4: Warp Scheduling

scratchpad memory (*Non-owner thread blocks*). We refer the to the warps in these thread blocks as 'Unshared Warps', 'Owner warps', and 'Non-owner Warps' respectively. When an owner thread block finishes its execution, it transfers its ownership to its corresponding non-owner thread block, and the new thread block that will be launched becomes the non-owner thread block.

Scheduling these warps in the SM plays an important role in improving the performance of applications. Hence we propose an optimization, *Owner Warp First (OWF)*, that schedules the warps in the following order: (1) Owner warps, (2) Unshared warps, and (3) Non-owner warps. Giving the highest priority to owner warps helps them in finishing sooner so that the dependent non-owner warps can resume their execution, which can help in hiding long execution latencies. Figure 4 shows the benefit of giving the first priority to owner warp when compared to non-owner warp. Consider an SM that has three warps: Unshared (U), Owner (O), and Non-owner (N) warps. Assume that they need to execute 3 instructions I_1, I_2, and I_3 as shown in the figure; the latency of *Add* and *Mov* instruction is 1 cycle, and the latency of *Load* instruction is 5 Cycles. Also, assume that instruction I_2 uses a shared scratchpad memory location. If unshared warp is given highest priority (shown as Unshared Warp First in the figure), then it can issue I_1 in the 1st cycle and issue I_2 in the 2nd cycle. However, it can not issue I_3 in the 3rd cycle since I_3 is dependent on I_2 for register R_2, and I_2 takes five cycles to complete its execution. However, the owner warp can execute I_1 in the 3rd cycle. Similarly, it can issue I_2 in the fourth cycle. The non-owner with least priority can start issuing I_1 in the 5th cycle, however, it can not issue I_2 in the 6th cycle since I_2 uses a shared scratchpad memory location, and it can access the shared location only after the owner thread block releases the lock, hence it waits until the owner warp finishes execution. Once the owner warp finishes the execution of I_2 and I_3 in the 8th and 9th cycles respectively, the non-owner resumes the execution of I_2 in 10th cycle, and it can subsequently finish in 15 cycles.

If owner warp is given first priority compared to unshared warp, it can issue I_1 and I_2 in its 1st and 2nd cycles respectively. Similarly, the unshared warp, with second priority, can issue I_1 and I_2 in 3rd and 4th cycles. The non-owner warp with least priority can issue I_1 in 5th cycle, and it waits for owner warp to release the shared scratchpad memory. Once the owner warp completes the execution of I_2 and I_3in 7th and 8th cycles, the non-owner can resume the execution by overlapping the execution of I_2 in the 8th cycle with unshared warp. Finally, the unshared warp and non-owner warp can finish their execution in 9th and 13th cycle respectively. Thus improving the overall performance.

4.1. Effect of Barriers

Our approach to scratchpad sharing does not require any special handling of barriers (`__syncthreads()`). Recall that in our approach, the scratchpad is shared at the thread block granularity. An owner thread block gets the lock on the shared scratchpad and releases only after it completes its execution. The warps from the non-owner thread block wait for its owner thread block for shared scratchpad memory. However, the warps from the owner thread block never wait for the warps of the non-owner thread block since they already have lock. In the presence of `__syncthreads()` instruction only the warps within the same thread blocks wait for other warps, and they make progress after all the warps of the thread block arrive at the barrier [CUDA 2012]. Hence a circular wait is not possible even in the presence of both barriers as well as shared scratchpad locks, thus avoiding deadlocks.

4.2. Scratchpad Sharing on Multiple SMs

GPGPUSIM [2014] uses a round robin scheduling algorithm for scheduling thread blocks on multiple SMs[2]. Consider a GPU that has p SMs. Assume that a kernel needs to launch N thread blocks ($B_1, B_2, \ldots, B_N$) in the GPU. Further, assume that default scratchpad allocation mechanism can launch m thread blocks in each SM. Thus, the i^{th} SM initially has thread blocks with ids $B_i, B_{p+i}, B_{2p+i}, \ldots B_{(m-1)p+i}$ in it. Whenever a SM finishes the execution of a thread block, a new thread block is launched in it, until all the N blocks are finished.

Our proposed scratchpad sharing mechanism does not modify the thread block scheduling mechanism. Assuming our approach launches n blocks ($n \geq m$), the i^{th} SM gets thread blocks with ids $B_i, B_{p+i}, B_{2p+i}, \ldots, B_{(n-1)p+i}$. The additional $n - m$ thread blocks (i.e., $B_{mp+i}, B_{(m+1)p+i}, \ldots$) share scratchpad memory with blocks $B_i, B_{p+i}, \ldots$ respectively. The remaining thread blocks remain in the unsharing mode. Note that the thread blocks that share scratchpad memory are always part of the the same SM. Whenever a new thread block is launched in place of an old block that has finished its execution, the new block gets the same scratchpad sharing status (shared/unshared) as the old one.

Since thread blocks in an SM can complete their execution in any order [CUDA 2012] (no priority among thread blocks), and the scratchpad memory is shared only among the thread blocks within a SM, priority inversion problem does not arise with the scratchpad sharing approach.

5. THE NEED FOR COMPILER OPTIMIZATIONS

In scratchpad sharing, when two thread blocks (say, TB_0 and TB_1) are launched in shared mode, one of them accesses the shared scratchpad region at a time. As soon as one thread block, say TB_0, starts accessing the shared scratchpad region, the other thread block, TB_1, can not access the shared scratchpad region and hence may have wait until TB_0 finishes execution.

Example 5.1. Consider the CFG in Figure 5, which is obtained for *SRAD1* benchmark application (Table I). In the figure, the program point marked **L** corresponds to the last access to the shared scratchpad. Without compiler assistance, the shared scratchpad region can be released only at the end of the last basic block (Exit node of CFG) even though it is never accessed after L. □

To promote the release of shared scratchpad region before the end of kernel execution, we introduce a new hardware instruction (PTX instruction) called *relssp*. Our proposed

[2]Note that this is different from NVIDIA GPU that is believed to use a FIFO policy [Pai et al. 2014].

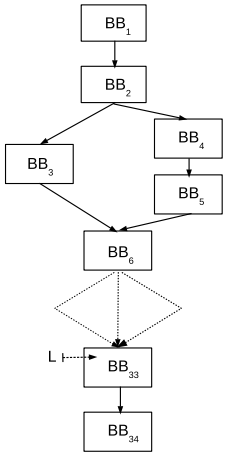

Fig. 5: Release of Shared Scratchpad

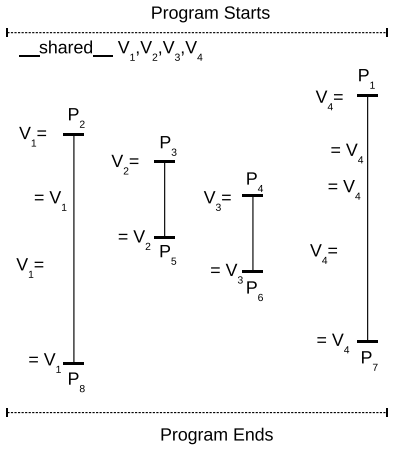

Fig. 6: Access Ranges of Scratchpad Variables

compiler optimization can place the *relssp* instructions in a kernel such that shared scratchpad memory is released as early as possible by each thread.

Example 5.2. Consider the scenario in Figure 6, where a kernel function declares four equal sized scratchpad variables V_1 to V_4. The figure also shows the regions of the kernel within which different variables are accessed. If V_1 and V_4 are allocated into shared scratchpad region, then the shared scratchpad region is accessed from program point P_1 to program point P_8. However, when V_2 and V_3 are allocated to shared scratchpad region, the shared region is accessed for a shorter duration, i.e., from program point P_3 to program point P_6. □

Note that the choice of allocation of scratchpad variables into shared and unshared scratchpad regions does not affect the correctness of the program but can affect the availability of the shared scratchpad region, and hence the effectiveness of sharing.

6. COMPILER OPTIMIZATIONS

In this section we describe a compile time memory allocation scheme and an analysis to optimally place *relssp* instructions. The memory allocation scheme allocates scratchpad variables into shared and unshared region such that shared scratchpad variables are accessed only for a small duration during the run-time. In the presence of loops where the number of iterations of loops are not computable at compile time, it is not possible to statically bound the number of instructions executed at run-time. Hence we need to use approximate loop bounds. Any approximation is fine since, as noted earlier, it only affects the effectiveness of sharing, but not the correctness[3].

To simplify the description of the required analyses, we make the following assumptions:

— The control flow graph (CFG) for a function (kernel) has a unique Entry and a unique Exit node.
— There are no *critical* edges in the CFG. A critical edge is an edge whose source node has more than one successor and the destination node has more than one predecessor.

[3]Profiling and user annotations can help in finding better approximations for the loop bounds. However we have not used these in our current implementation.

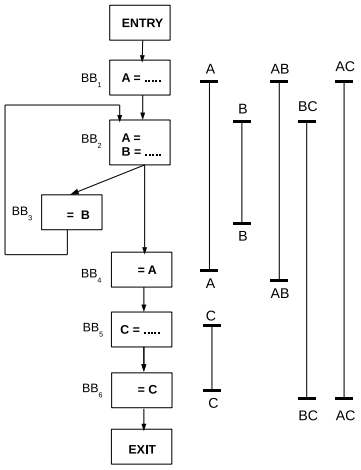

Fig. 7: Access Ranges of Variables

Table III: Access Ranges for Scratchpad Variables and Sets.
t denotes true, **f** denotes false. Sets of variables are written as concatenation of variables. For example, AB denotes {A, B}.

	For Variables						For Sets of Variables					
	IN			OUT			IN			OUT		
Block	A	B	C	A	B	C	AB	BC	CA	AB	BC	CA
Entry	f	f	f	f	f	f	f	f	f	f	f	f
BB1	f	f	f	t	f	f	f	f	f	t	f	t
BB2	t	t	f	t	t	f	t	t	t	t	t	t
BB3	t	t	f	t	t	f	t	t	t	t	t	t
BB4	t	f	f	f	f	f	t	t	t	f	t	t
BB5	f	f	f	f	f	t	f	t	t	f	t	t
BB6	f	f	t	f	f	f	f	t	t	f	f	f
Exit	f	f	f	f	f	f	f	f	f	f	f	f

These assumptions are not restrictive as any control flow graph can be converted to the desired form using a preprocessing step involving simple graph transformations: adding a source node, adding a sink node, and adding a node to split an edge [Khedker et al. 2009; Kam and Ullman 1976; Muchnick 1997].

6.1. Minimizing Shared Scratchpad Region

Consider a GPU that uses scratchpad sharing approach such that two thread blocks involved in sharing can share a fraction $f < 1$ of scratchpad memory. Assume that each SM in the GPU has M bytes of scratchpad memory, the kernel that is to be launched into the SM has N scratchpad variables, and each thread block of the kernel requires M_{tb} bytes of scratchpad memory. We allocate a subset S of scratchpad variables into shared scratchpad region such that: (1) The total size of the scratchpad variables in the set S is equal to the size of shared scratchpad ($f \times M_{tb}$), and (2) The region of access for variables in S is minimal in terms of the number of instructions.

To compute the region of access for S, we define *access range* for a variable as follows:

DEFINITION 6.1. **Access Range of a Variable:** *A program point π is in the* access range *of a variable v if both the following conditions hold: (1) There is an access (definition or use) of v on some path from* Entry *to π* and *(2) There is an access of v on some path from π to* Exit.

Intuitively, the access range of a variable covers every program point between the first access and the last access of the variable in an execution path. The access range for a variable can contain disjoint regions due to branches in the flow graph.

DEFINITION 6.2. **Access Range of a Set of Variables:** *A program point π is in the* access range *of a set of variable S if both the following conditions hold: (1) There is an access to a variable $v \in S$ on a path from* Entry *to π* and *(2) There is an access to a variable $v' \in S$ on a path from π to* Exit.

Example 6.1. Consider a kernel whose CFG is shown in Figure 7. The kernel uses 3 scratchpad variables A, B and C. Variable A is accessed in the region from basic block BB_1 to basic block BB_4. The start of basic block BB_2 is considered in access range of A because there is a path from Entry to start of BB_2 that contains an access of A

(definition in BB_1) and there is a path from the start of BB_2 to Exit that contains the access of A (use in BB_4).

Consider the set $S = \{B, C\}$. Basic block BB_4 is in access range of S because there is a path from Entry to BB_4 containing the access of B (definition in BB_2), and there is a path from BB_4 to Exit containing the access of the C (use in BB_6). □

To compute the access ranges for a program, we need a forward analysis to find the first access of the scratchpad variables, and a backward analysis to find the last access of the scratchpad variables. We define these analyses formally using the following notations:

— $\mathsf{IN}(BB)$ denotes the program point before the first statement of the basic block BB. $\mathsf{OUT}(BB)$ denotes the program point after the last statement of BB.
— $\mathsf{PRED}(BB)$ denotes the set of predecessors, and $\mathsf{SUCC}(BB)$ denotes the set of successors of BB.
— $\mathsf{PreIN}(v, BB)$ is true if there is an access to variable v before $\mathsf{IN}(BB)$. $\mathsf{PreOUT}(v, BB)$ is true if there is an access to the variable v before $\mathsf{OUT}(BB)$.
— $\mathsf{PostIN}(v, BB)$ is true if there is an access to variable v after $\mathsf{IN}(BB)$. $\mathsf{PostOUT}(v, BB)$ is true if there is a access to variable v after $\mathsf{OUT}(BB)$.
— $\mathsf{AccIN}(S, BB)$ is true if $\mathsf{IN}(BB)$ is in access range of a set of scratchpad variables S. $\mathsf{AccOUT}(S, BB)$ is true if $\mathsf{OUT}(BB)$ is in access range of a set of scratchpad variables S.

The data flow equations to compute the information are[4]:

$$\mathsf{PreOUT}(v, BB) = \begin{cases} \text{true, if } BB \text{ has an access of } v \\ \mathsf{PreIN}(v, BB), \text{ otherwise} \end{cases} \qquad \mathsf{PreIN}(v, BB) = \begin{cases} \text{false, if } BB \text{ is Entry block} \\ \bigvee\limits_{BP \in \mathsf{PRED}(BB)} \mathsf{PreOUT}(v, BP), \text{ otherwise} \end{cases}$$

$$\mathsf{PostIN}(v, BB) = \begin{cases} \text{true, if } BB \text{ has an access of } v \\ \mathsf{PostOUT}(v, BB), \text{ otherwise} \end{cases} \qquad \mathsf{PostOUT}(v, BB) = \begin{cases} \text{false, if } BB \text{ is Exit block} \\ \bigvee\limits_{BS \in \mathsf{SUCC}(BB)} \mathsf{PostIN}(v, BS), \text{ otherwise} \end{cases}$$

We decide whether the access range of a set of scratchpad variables S includes the points $\mathsf{IN}(BB)$ and $\mathsf{OUT}(BB)$ as:

$$\mathsf{AccIN}(S, BB) = \left(\bigvee_{v \in S} \mathsf{PreIN}(v, BB)\right) \bigwedge \left(\bigvee_{v \in S} \mathsf{PostIN}(v, BB)\right)$$
$$\mathsf{AccOUT}(S, BB) = \left(\bigvee_{v \in S} \mathsf{PreOUT}(v, BB)\right) \bigwedge \left(\bigvee_{v \in S} \mathsf{PostOUT}(v, BB)\right)$$

Example 6.2. Table III shows the program points in the access ranges of scratchpad variables for CFG of Figure 7. The table also shows the program points in the access ranges of sets of two scratchpad variables each. □

Let SV denote the set of all scratchpad variables. For every subset S of SV having a total size equal to the size of shared scratchpad memory, our analysis counts the total number of instructions in the access range of S. Finally the subset that has the minimum count is selected for allocation in the shared scratchpad memory.

Example 6.3. Consider once again the CFG in Figure 7. For simplicity, assume that all the variables have equal sizes, and each basic block contains the same number of instructions. Consider a scratchpad sharing approach that can allocate only two of the variables into the shared scratchpad region. From the CFG, and from Table III, it is clear that when A and B are allocated into shared scratchpad memory, the shared

[4]The analysis can be extended easily to compute information at any point inside a basic block. We ignore it for brevity.

```
1 void ReleaseSSP()
2 {
3     static int count=0;
4     ++count;
5     if(count==ACTIVE_THREADS){
6         count=0;
7         UNLOCKSHAREDREGION();
8     }
9 }
```

Fig. 8: Pseudocode of *relssp* Instruction

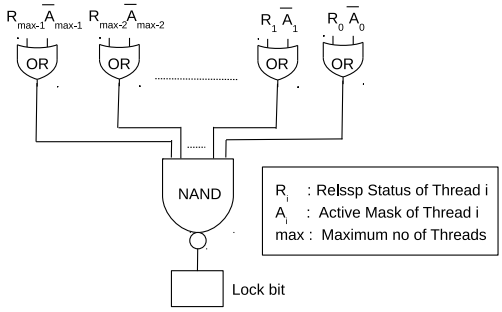

Fig. 9: Hardware Implementation of *relssp* Instruction

region is smaller, compared to when either {B,C} or {A,C} are allocated in the shared region. □

6.2. Implementation of *relssp* Instruction

In scratchpad sharing approach, a shared thread block acquires a lock before accessing shared scratchpad region and unlocks it only after finishing its execution. This causes a delay in releasing the shared scratchpad because the thread block holds the scratchpad memory till the end of its execution, even though it has finished accessing shared region.

To minimize the delay in releasing the shared scratchpad, we propose a new instruction, called *relssp*, in PTX assembly language. The semantics of *relssp* instruction is to unlock the shared region only when all active threads within a thread block finished executing the shared region. Figure 8 shows the pseudo code for *relssp* instruction. The RELEASESSP() procedure maintains **count**, an integer initialized to zero. When an active thread within a thread block executes a *relssp* instruction, it increments the **count** value. When all active threads of a thread block execute *relssp* instruction (Line 5, when **count** equals ACTIVE_THREADS), the shared region is unlocked by invoking UNLOCKSHAREDREGION(). The unlock procedure releases the shared scratchpad region by resetting the lock variable. The execution of *relssp* by a thread block that does not access shared scratchpad region has no effect.

It is clear that **count** in Figure 8 has to be a shared variable, hence a software implementation will require to manage critical section. The same algorithm, however, can be efficiently implemented in hardware circuit as shown in Figure 9. The i^{th} thread within a thread block is associated with an active mask (A_i) and a release bit (R_i). The mask A_i is set if the i^{th} thread is active. When this thread executes *relssp* instruction, the release bit (R_i) gets set. The shared scratchpad region is unlocked only when all the active threads in a thread block execute *relssp* instruction (the lock bit, i. e. the output of NAND gate becomes 0 in Figure 9). In other words, shared scratchpad region is unlocked if $\forall i \; A_i \rightarrow R_i$ is true.

6.3. Algorithm for optimal placement of *relssp* instruction

In Section 6.2, we introduced a new instruction to release the shared scratchpad memory. In this section, we discuss a compile-time analysis for optimal insertion of *relssp* instruction in the program. We insert a *relssp* instruction at a program point π such that the following conditions are met:

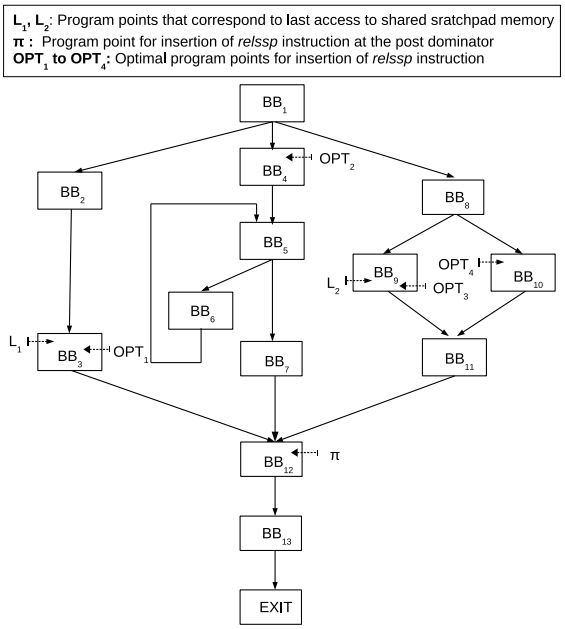

Fig. 10: Possible Insertion Points for *relssp*

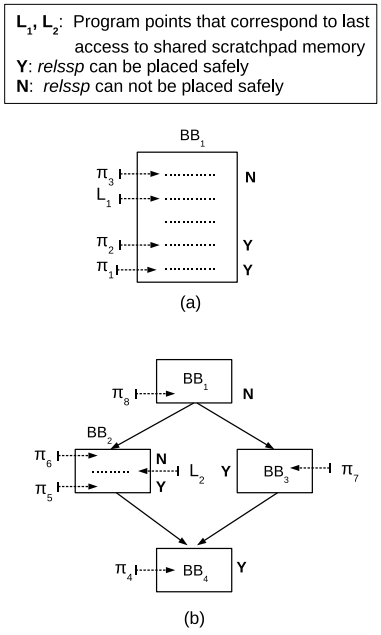

Fig. 11: Scenarios for Optimal Insertion of *relssp*

CONDITION 1. **Safety:** *The relssp instruction must be executed by each active thread within a thread block, and it must be executed after last access to shared scratchpad memory.*

CONDITION 2. **Optimality:** *The relssp instruction must be executed by each active thread exactly once.*

Condition 1 ensures that shared scratchpad is eventually released by a thread block since the instruction is executed by all the threads of a thread block. Also, it guarantees that shared scratchpad is released only after a thread block has completed using it. Whereas, Condition 2 avoids redundant execution of *relssp* instruction.

In the scratchpad sharing, a thread block releases the shared scratchpad memory after completing its execution, hence it is equivalent to having a *relssp* instruction placed at the end of the program, which guarantees both the conditions, albeit at the cost of delay in releasing the shared scratchpad. A simple improvement that promotes early release of shared scratchpad memory and ensures both the conditions, is to place the *relssp* instruction at a basic block $BB_{postdom}$ where $BB_{postdom}$ is a common post dominator of those basic blocks having the last accesses to the shared scratchpad memory along different paths. Further, $BB_{postdom}$ should dominate Exit, i.e., it should be executed in all possible execution paths. As the following example shows, this strategy, though an improvement over placing *relssp* in Exit, may also result in delaying the release of shared scratchpad memory.

Example 6.4. Consider a CFG shown in Figure 10. Assume that L_1, L_2 denote the program points that correspond to the last accesses to shared scratchpad memory. Since *relssp* instruction is to be executed by all the threads of thread block, it can be placed at the post dominator of the basic blocks BB_3 and BB_9, i.e., program point marked π in BB_{12}, which is visible to all threads. However, this delays the release of shared scratchpad. Consider a thread that takes a path along BB_9, it can execute *relssp* immediately after executing the last access to shared region (shown as OPT_3

in the figure). It executes *relssp* at program point π. Similarly, when a thread takes a path along the basic block BB_4, it releases the shared scratchpad at π even though it does not access any shared scratchpad in that path. The scratchpad can be released at program point OPT_2 in BB_4. □

As is clear from the above example, placement of *relssp* instruction has an effect on the availability of shared scratchpad memory. Intuitively, a safely placed *relssp* instruction at a program point π can be moved to a previous program point π' in the same basic block provided the intervening instructions do not access shared scratchpad. The movement of *relssp* from a basic block BB to predecessor BB' is possible provided every other successor of BB' also does so.

Example 6.5. Figure 11(a) shows a basic block BB_1, which has the last access to the shared scratchpad memory at L_1. In this block, if the *relssp* instruction can be placed safely at the program point π_1, then it can be moved to π_2 since there is no access to shared scratchpad memory between π_1 and π_2. However, it can not be moved to the program point π_3 within the same basic block, because it violates safety (Condition 1).

Consider another scenario shown in Figure 11(b), basic block BB_2 has the last access to shared memory at L_2, and basic blocks BB_1, BB_3, and BB_4 do not access any scratchpad memory. If the *relssp* instruction can be placed safely at π_4 in BB_4, then it can be moved to a program point π_5 and π_7 in the basic blocks BB_2 and BB_3 respectively. However, it can not be moved to program point π_6 in BB_2 and π_8 in BB_1 since it violates of Condition 1. Also, the *relssp* instruction can not be moved from π_7 in BB_3 to π_8 in BB_1 since the basic block BB_2, which is a successor of BB_1, does not allow the *relssp* instruction to be placed at π_8. □

We now formalize these intuitions into a backward data flow analysis. The notations used are:

— $\mathsf{IN}(BB)$ denotes the program point before the first statement of the basic block BB. $\mathsf{OUT}(BB)$ denotes the program point after the last statement of BB.
— $\mathsf{SafeIN}(BB)$ is true if the *relssp* instruction can be safely placed at $\mathsf{IN}(BB)$, and $\mathsf{SafeOUT}(BB)$ is true if the *relssp* instruction can be safely placed at $\mathsf{OUT}(BB)$.
— INS_π, if true, denotes that *relssp* will be placed at program point π by the analysis.

The data flow equations are:

$$\mathsf{SafeIN}(BB) = \begin{cases} \text{false, if } BB \text{ has shared} \\ \qquad \text{scratchpad access} \\ \mathsf{SafeOUT}(BB), \text{ otherwise} \end{cases} \qquad \mathsf{SafeOUT}(BB) = \begin{cases} \text{true, if } BB \text{ is Exit block} \\ \bigwedge\limits_{BS \in \mathsf{SUCC}(BB)} \mathsf{SafeIN}(BS), \text{ otherwise} \end{cases}$$

The above equations compute the program points where *relssp* can be placed safely. For a basic block BB, $\mathsf{OUT}(BB)$ is an optimal place for *relssp* instruction, if *relssp* can be placed safely at $\mathsf{OUT}(BB)$, and it can not be moved safely to its previous program point in the basic block, i.e., $\mathsf{IN}(BB)$ is false. This is computed as:

$$\mathsf{INS}_{\mathsf{OUT}(BB)} = \mathsf{SafeOUT}(BB) \wedge \neg(\mathsf{SafeIN}(BB)) \tag{1}$$

Similarly, $\mathsf{IN}(BB)$ is an optimal point for *relssp* instruction, when the instruction can not be moved to its predecessors [5]. This can be computed as:

$$\mathsf{INS}_{\mathsf{IN}(BB)} = \neg \left(\bigwedge_{BP \in \mathsf{PRED}(BB)} \mathsf{SafeOUT}(BP) \right) \wedge \mathsf{SafeIN}(BB) \tag{2}$$

Equations (1) and (2) together, along with the absence of critical edges, ensure the optimality condition that each thread executes the *relssp* instruction exactly once.

[5] Absence of critical edges guarantees that the instruction can either be moved to *all* predecessors or to *none*.

7. REQUIREMENTS FOR SCRATCHPAD SHARING

7.1. Hardware Requirements

Figure 12 shows the modified GPU architecture to implement scratchpad sharing. Our approach requires two modifications to scheduler unit in the SM. The first change is to the warp scheduler which uses OWF optimization. The second change is the inclusion of resource access unit, which follows the scratchpad access mechanism as discussed in Section 3. The resource access unit requires the following additional storage:

(1) Each SM requires a bit (shown as *ShSM* in Figure 12) to indicate whether the scratchpad sharing is enabled for it. This bit is set when the number of thread blocks launched using our approach is more than that of baseline approach.
(2) Every thread block involved in sharing stores the id of its partner thread block in the *ShTB* table. If a thread block is in unsharing mode, a -1 is stored. For T thread blocks in the SM, we need a total of $T \log_2(T+1)$ bits.
(3) Each warp a requires a bit to specify if it is an owner warp. For W warps in the SM, W bits are needed.
(4) For each pair of shared thread blocks, a lock variable is needed in order to access shared scratchpad memory. This variable is set to the id the thread block that acquired shared scratchpad memory. For T thread blocks, there are at most $\lfloor T/2 \rfloor$ pairs of sharing thread blocks in the SM. This requires $\lfloor \frac{T}{2} \rfloor \lceil (\log_2 T) \rceil$ bits in the SM.

If a GPU has N SMs and allows a maximum of T thread blocks and W warps per SM, then the number of additional bits required is: $\big(1 + T\log_2(T+1) + W + \lfloor \frac{T}{2} \rfloor \lceil (\log_2 T) \rceil\big) * N$. For the architecture we used for simulation (shown in Table II), the overhead is 209 bits per SM. In addition, each scheduler unit in the SM requires two comparator circuits and one arithmetic circuit to set the lock (See Figure 3).

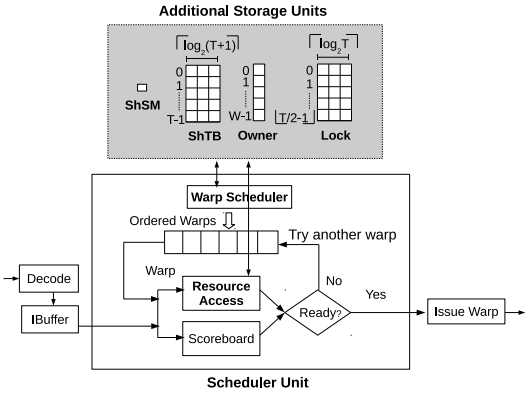

Fig. 12: Modified Architecture for Scratchpad Sharing

7.2. Analysis of Compiler Optimizations

The dataflow analyses to compute definitions and usages of scratchpad variables (Section 6.1) are bit-vector data flow analyses [Khedker et al. 2009]. For a kernel with n scratchpad variables and m nodes (basic blocks) in the flow graph, the worst case complexity is $O(n \times m^2)$ (assuming set operations on n bit-wide vectors take $O(n)$ time).

The computation of access ranges for sets of variables may require analyzing all $O(2^n)$ subsets in the worst case, where the largest size of a subset is $O(n)$. Thus, given the usage and definitions at each program point in the kernel, computation of AccIN and AccIN requires $O(m \times n \times 2^n)$ time. Therefore, the total time complexity is $O(n \times m^2 + m \times n \times 2^n)$. Since the number of scratchpad variables in a kernel function is small (typically, $n \leq 10$), the overhead of the analysis is practical.

Our approach inserts *relssp* instructions in a CFG such that *relssp* is called exactly once along any execution path. In the worst case, all nodes in a CFG (except Entry and Exit blocks) might fall along different paths from Entry to Exit. Hence the worst case number of *relssp* inserted is $O(m)$.

Table IV: Set-3 Benchmarks: The Number of Thread Blocks is Not Limited by Scratchpad Memory

Benchmark	Application	Kernel	Limited by
GPGPU-SIM	BFS	Kernel	Threads, Registers
RODINIA	b+tree	findRangeK	Registers
CUDA-SDK	dct8x8_5 (DCT5)	CUDAkernel1DCT	Blocks
RODINIA	gaussian	FAN1	Threads
GPGPU-SIM	NN	executeSecondLayer	Blocks

Table V: Details of Modifications to the Benchmarks

Application	File Name	Line Number	Constant/Variable Name	Old Value	New Value
backprop	backprop.h	10	WIDTH	16	48
		11	HEIGHT	16	48
DCT[1..4]	dct8x8.cu	317	numIterations	100	1
NQU	nqueen.c	11	THREAD_NUM	96	64
SRAD1 SRAD2	srad.h	3	BLOCK_SIZE	16	24
heartwall	define.c	9	NUMBER_THREADS	512	128
		13	ALL_POINTS	51	140
	main.cu [6]	96	frame_processed	5	1
histogram	histogram_common.h	18	HISTOGRAM64_BIN_COUNT	256	384
NW1 NW2	needle.h	1	BLOCK_SIZE	32	16

8. EXPERIMENTAL EVALUATION

We implemented the proposed scratchpad sharing approach and integrated *relssp* instruction in GPGPU-Sim V3.x [GPGPUSIM 2014] simulator. We implemented the compiler optimizations in PTX assembly [PTX 2014] using Ocelot [Diamos et al. 2010] framework. The baseline architecture that we used for comparing our approach is shown in Table II. We evaluated our approach on several kernels from CUDA-SDK [CUDA-SDK 2014], GPGPU-Sim [Bakhoda et al. 2009], and Rodinia [Che et al. 2009] benchmark suites.

Depending on the amount and the last usage of the shared scratchpad memory by the applications, we divided the benchmark applications into three sets. Set-1 and Set-2 (Table I) consists of applications whose number of resident thread blocks are limited by scratchpad memory. For Set-1, the applications do not access scratchpad memory till towards the end of their execution, while for Set-2, the applications access scratchpad memory till towards the end of their execution. The introduction of *relssp* instruction is expected to give benefit over our earlier approach [Jatala et al. 2016] only for Set-1 applications. Set-3 benchmarks (Table IV) consist of applications whose number of thread blocks are not limited by scratchpad memory, but by some other parameter. These are included to show that our approach does not negatively affect the performance of applications that are not limited by scratchpad memory.

For each application in Set-1 and Set-2 benchmarks, Table I shows the kernel that is used for evaluation, the number of the scratchpad variables declared in each kernel, the amount of the scratchpad memory required for each thread block, and the thread block size. Some applications in Set-1 and Set-2 benchmarks are modified to make sure that the number of thread blocks is limited by scratchpad memory, thus making scratchpad sharing approach applicable. These changes increase the scratchpad memory requirement per thread block and are shown in Table V. For Set-3 benchmarks, Table IV shows the cause of limitation on the number of thread blocks. The causes

[6]This has been changed for quick simulation.

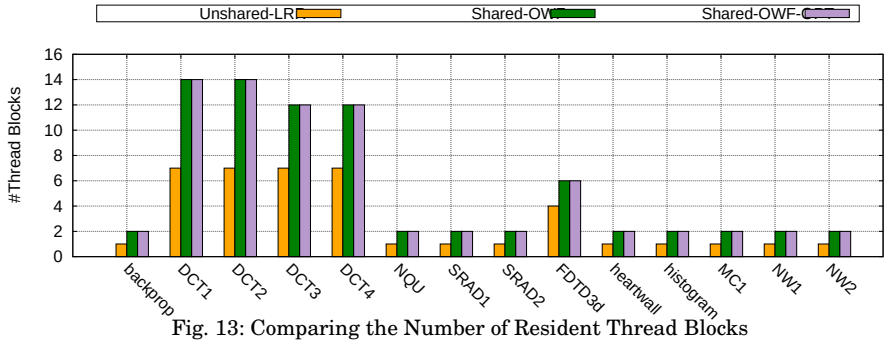

Fig. 13: Comparing the Number of Resident Thread Blocks

include the limit on the number of registers, the maximum limit on the number of resident thread blocks, and the maximum limit on the number of resident threads.

We compiled all the applications using CUDA 4.0[7] and simulated them using the GPGPU-Sim simulator. We use a threshold (t) to configure the amount of scratchpad sharing. If each thread block requires R_{tb} amount of scratchpad memory, then we allocate $R_{tb}(1+t)$ for each pair of shared thread blocks, in which we allocate $R_{tb}(1-t)$ as shared scratchpad memory. We analyzed the benchmark applications for various threshold values and choose the value t as 0.1 (i.e., 90% scratchpad is shared among pair of thread blocks) to give the maximum benefit. The details of the experiments to choose t are given in technical report [Jatala et al. 2015].

We measure the performance of our approach using the following metrics:

(1) The **number of the resident thread blocks** launched in the SMs. This is a measure of the amount of thread level parallelism present in the SMs.
(2) The **number of instructions executed per shader core clock cycle (IPC)**. This is a measure of the throughput of the GPU architecture.
(3) The **number of simulation cycles** that an application takes to complete its execution. This is a measure of the performance of the benchmark applications.

8.1. Analysis of Set-1 and Set-2 Benchmarks

We use *Unshared-LRR* to denote the baseline unsharing approach, *Shared-OWF* to denote our scratchpad sharing approach with OWF scheduler, and *Shared-OWF-OPT* to denote the scratchpad sharing approach that includes OWF scheduler and compiler optimizations.

8.1.1. Comparing the Number of Resident Thread Blocks. Figure 13 shows the number of thread blocks for the three approaches. For applications *DCT1* and *DCT2*, *Unshared-LRR* launches 7 thread blocks in the SM according to the amount of scratchpad memory required by their thread blocks. *Shared-OWF* launches 14 thread blocks in the SM, where each of the 7 additional thread blocks share scratchpad memory with other resident thread blocks. For *DCT3* and *DCT4* applications, *Unshared-LRR* launches 7 thread blocks in the SM, whereas *Shared-OWF* launches 12 thread blocks in the SM such that the additional 5 thread blocks share scratchpad memory with the existing 5 thread blocks; while the remaining 2 existing thread blocks in the SM do not share scratchpad memory with any other thread block. For *FDTD3d*, *Shared-OWF* launches 2 additional thread blocks in the SM when compared to *Unshared-LRR*, which share scratchpad memory with other 2 resident thread blocks. For the remaining applications, *Unshared-LRR* launches 1 thread block, whereas *Shared-OWF* launches 1 ad-

[7]GPGPU-Sim and Ocelot do not support CUDA 5.0 and above.

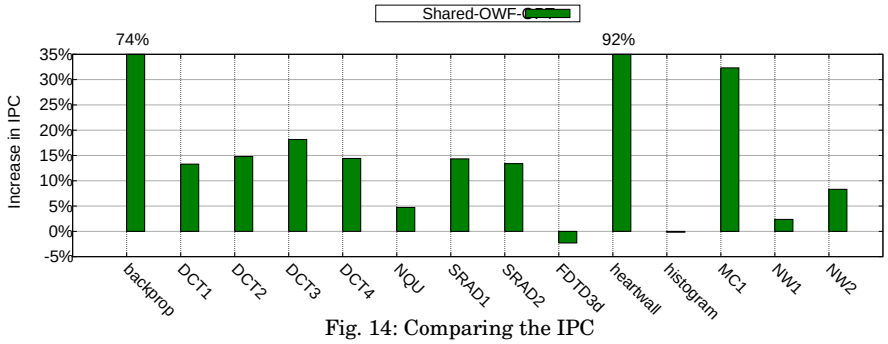

Fig. 14: Comparing the IPC

ditional thread block in the SM which shares scratchpad memory with the existing thread block. Note that the number of thread blocks launched by *Shared-OWF-OPT* is exactly same as that of *Shared-OWF*. This is expected since the number of additional thread blocks launched by scratchpad sharing approach depends on two parameters: (1) the amount of scratchpad sharing, and (2) the amount of scratchpad memory required by a thread block; and our compiler optimizations do not affect either of these parameters.

8.1.2. Performance Comparison. Figure 14 compares the performance of *Shared-OWF-OPT*[8] in terms of the number of instructions executed per cycle (IPC) with that of *Unshared-LRR*. We observe a maximum improvement of 92.17% and an average (Geometric Mean) improvement of 19% with *Shared-OWF-OPT*. The maximum benefit of 92.17% is for *heartwall* because the additional thread blocks launched by *Shared-OWF-OPT* do not access the shared scratchpad region. Hence all the additional thread blocks make progress without waiting for corresponding shared thread blocks. *MC1* improves by 32.32% because additional thread blocks launched in the SM make significant progress before accessing shared scratchpad region. *backprop* shows an improvement of 74.2%, it leverages both scratchpad sharing and the compiler optimizations to perform better. The improvements in *SRAD1* and *SRAD2* applications are largely due to the compiler optimizations. *FDTD3d* slows down (–2.29%) with *Shared-OWF-OPT* due to more number of L1 and L2 cache misses when compared to *Unshared-LRR*. *histogram* does not benefit from sharing since the thread blocks start accessing shared scratchpad region early in the execution, causing one of the blocks from each sharing pair to wait for the lock.

8.1.3. Overhead of relssp Instruction. Table VI shows the run-time overhead of inserting *relssp* instruction. We report sum of the number of instructions executed by all threads for *Unshared-LRR*, *Shared-OWF*, and *Shared-OWF-OPT*. We also report the number of threads launched.

From the table, we observe that the number of instructions executed by *Unshared-LRR* and *Shared-OWF* is same. This is because *Shared-OWF* does not insert *relssp* instruction, and hence the input PTX assembly is not altered. *Shared-OWF-OPT* increases number of executed instructions as it inserts *relssp* and, in some cases, *GOTO* instruction to split critical edges. For the applications *DCT1*, *DCT2*, *SRAD1*, *SRAD2*, *NW1*, and *NW2*, the number of additionally executed instructions (shown as *Difference* (SO-U) in the table) is equal to number of threads because *Shared-OWF-OPT* inserts only the *relssp* instruction. Further, each thread executes *relssp* exactly once.

[8]IPC for *Shared-OWF-OPT* also takes into account the extra instructions inserted by the compiler optimizations. The absolute IPC values are shown in Table XIII in Appendix A.

Table VI: Comparing the Number of Simulated Instructions

Benchmark	Threads	Unshared-LRR (U)	Shared-OWF (S)	Shared-OWF-OPT (SO)	Difference (SO - U)
backprop	1,048,576	131,203,072	131,203,072	133,234,688	2,031,616
DCT1	32,768	9,371,648	9,371,648	9,404,416	32,768
DCT2	32,768	9,502,720	9,502,720	9,535,488	32,768
DCT3	32,768	11,255,808	11,255,808	11,304,960	49,152
DCT4	32,768	11,157,504	11,157,504	11,206,656	49,152
NQU	24,576	1,282,747	1,282,747	1,331,515	48,768
SRAD1	4,161,600	756,433,955	756,433,955	760,595,555	4,161,600
SRAD2	4,161,600	450,077,975	450,077,975	454,239,575	4,161,600
FDTD3d	144,384	5,549,531,392	5,549,531,392	5,549,820,160	288,768
heartwall	17,920	11,280,920	11,280,920	11,316,760	35,840
histogram	46,080	893,769,168	893,769,168	893,861,328	92,160
MC1	3,008	2,881,568	2,881,568	2,887,584	6,016
NW1	3,184	5,580,458	5,580,458	5,583,642	3,184
NW2	3,168	5,561,919	5,561,919	5,565,087	3,168

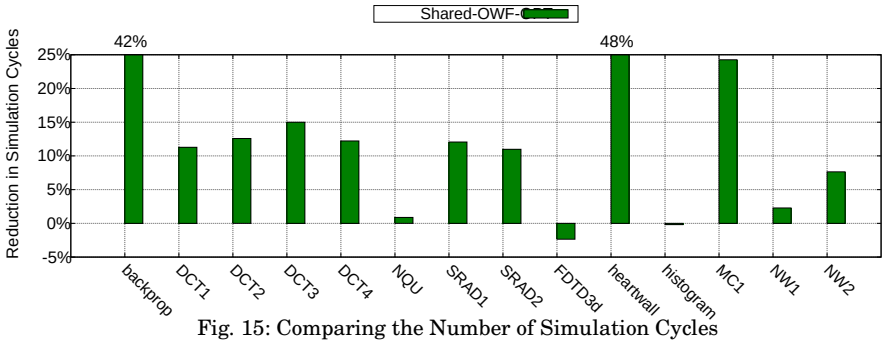

Fig. 15: Comparing the Number of Simulation Cycles

For *FDTD3d*, *heartwall*, *histogram*, and *MC1* applications, the number of additional instructions executed by *Shared-OWF-OPT* is twice that of number of threads. For these applications, each thread executes two additional instructions, i.e., one *relssp* instruction, and one *GOTO* instruction for splitting a critical edge. For *backprop*, *DCT3*, *DCT4*, and *NQU* applications, some threads take a path that has two additional instructions ($GOTO$ and *relssp*), while other threads take the path which has one additional *relssp* instruction.

8.1.4. Reduction in Simulation Cycles. Figure 15 shows the effectiveness of *Shared-OWF-OPT* by comparing the number of simulation cycles with that of *Unshared-LRR*. We observe a maximum reduction of 47.8% and an average reduction of 15.42% in the number of simulation cycles when compared to *Unshared-LRR*. Recall that *Shared-OWF-OPT* causes applications to execute more number of instructions (Table VI). These extra instructions are also counted while computing the simulation cycles for *Shared-OWF-OPT*.

8.1.5. Effectiveness of Optimizations. Figure 16 shows the effectiveness of our optimizations with scratchpad sharing. We observe that all applications, except *FDTD-3d* and *histogram*, show some benefit with scratchpad sharing even without any optimizations (shown as *Shared-NoOpt* in the figure). With OWF scheduling (*Shared-OWF*), applications improve further because OWF schedules the resident warps in a way that the non-owner warps help in hiding long execution latencies. For our benchmarks, minimizing shared scratchpad region (shown as *Shared-OWF-Reorder*) does not have any noticeable impact. This is because (a) Most applications declare only a single scratchpad variable (Table I) in their kernel, hence the optimization is not applicable (there is only one possible order of scratchpad variable declarations); and (b) For the remaining

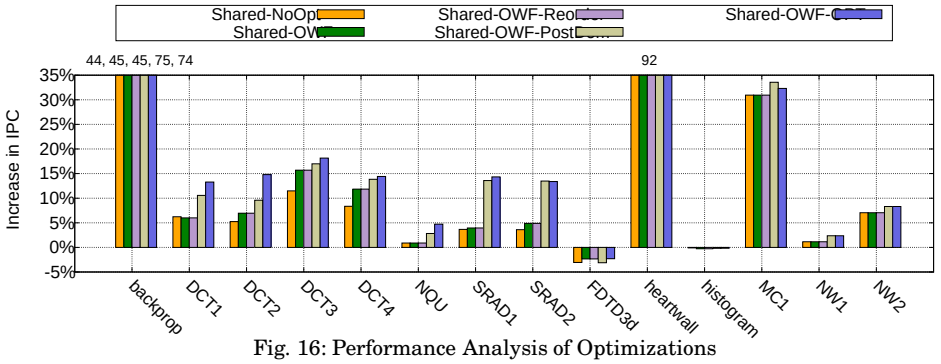

Fig. 16: Performance Analysis of Optimizations

applications, the scratchpad declarations are already ordered in the optimal fashion, i.e., the access to shared scratchpad region is already minimal.

The addition of *relssp* instruction at the postdominator and at the optimal places is denoted as *Shared-OWF-PostDom* and *Shared-OWF-OPT* respectively. All Set-1 applications improve with either of these optimizations because the *relssp* instruction helps in releasing the shared scratchpad memory earlier. For *backprop* and *SRAD2* applications, *Shared-OWF-PostDom* is better than *Shared-OWF-OPT* because the threads in *backprop* execute one additional *GOTO* instruction with *Shared-OWF-OPT* (*Shared-OWF-PostDom* does not require critical edge splitting). *SRAD2* has more number of stall cycles with *Shared-OWF-Opt* as compared to *Shared-OWF-PostDom*. For most of the other benchmarks, *Shared-OWF-Opt* performs better as it can push *relssp* instruction earlier than with *Shared-OWF-PostDom*, thus releasing shared scratchpad earlier allowing for more thread level parallelism.

As expected, Set-2 applications do not show much benefit with *Shared-OWF-PostDom* or *Shared-OWF-OPT* since they access shared scratchpad memory till towards the end of their execution. Hence both the optimizations insert *relssp* instruction in the Exit block in the CFGs. The application *heartwall* does not use shared scratchpad memory and hence it shows maximum benefit even without the insertion of *relssp* instruction.

8.1.6. Progress of Shared Thread Blocks. Figure 17 shows the effect of compiler optimizations by analyzing the progress of shared thread blocks through shared and unshared scratchpad regions. In the figure, *NoOpt* denotes the default scratchpad sharing mechanism where none of our optimizations are applied on an input kernel. *Minimize* denotes the scratchpad sharing approach which executes an input kernel having minimum access to shared scratchpad region. *PostDom* and *OPT* use our modified scratchpad sharing approach that execute an input kernel with additional *relssp* instructions placed at post dominator and optimal places (Section 6.3) respectively. In the figure, we show the percentage of simulation cycles spent in unshared scratchpad region (before acquiring shared scratchpad), shared scratchpad region, and unshared scratchpad region again (after releasing the shared scratchpad) respectively.

From Figure 17 we observe that shared thread blocks in all the applications access unshared scratchpad region before they start accessing shared scratchpad memory. Hence all the shared thread blocks can make some progress without wait. This progress is the main reason for the improvements seen with scratchpad sharing approach. Consider the application *heartwall*, where none of the shared thread blocks accesses shared scratchpad memory. Thus, all the shared thread blocks in the application spend their execution in the unshared scratchpad region. The compiler optimizations can not improve the progress of shared thread blocks any further. *Minimize* does

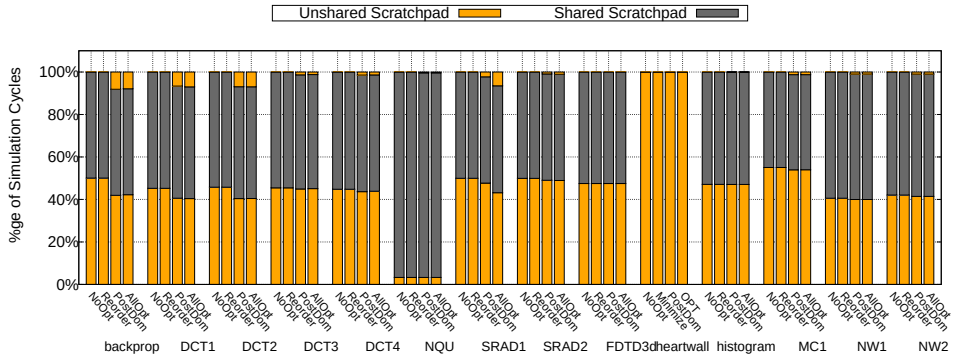

Fig. 17: Progress of Shared Thread Blocks

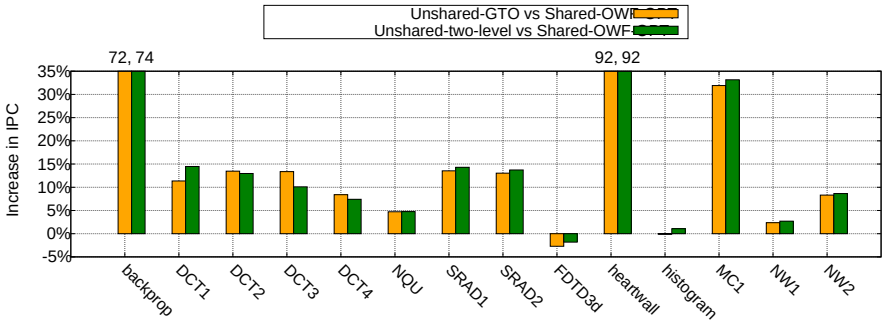

Fig. 18: Improvement in IPC for *Shared-OWF-OPT* w.r.t. baseline having (a) GTO scheduler, (b) two-level scheduler

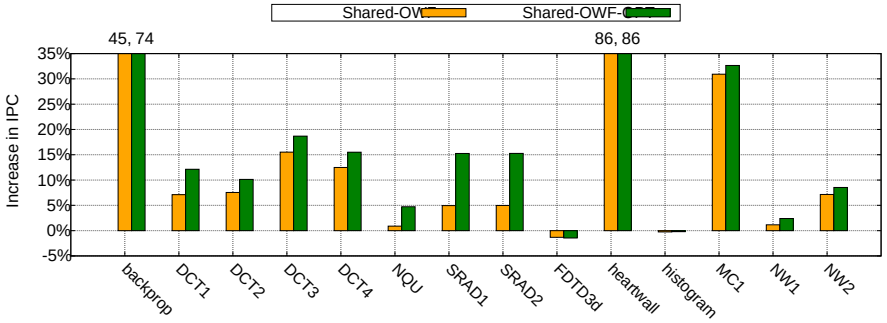

Fig. 19: Performance Analysis with GPU Configuration: Scratchpad=16K, L1-Cache=48K

not affect *DCT1*, *DCT2*, *DCT3*, *DCT4*, *FDTD3d*, *histogram* applications because the kernels in these applications declare single scratchpad variable. For the remaining applications, *Minimize* has same effect as that of *NoOpt*, because the default input PTX kernel already accesses the shared scratchpad variables such that access to shared scratchpad is minimum. We also observe that *PostDom* and *OPT* approaches improve only those applications that spend considerable simulation cycles in unshared scratchpad region *after* last access to shared scratchpad region.

8.1.7. Comparison with Different Schedulers. Figure 18 shows the effect of using different scheduling policies. The performance of *Shared-OWF-OPT* approach is compared with the baseline unshared implementation that uses greedy then old (GTO) and two-level

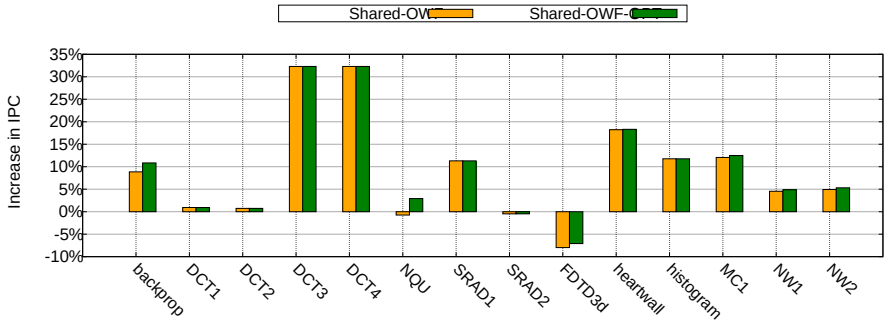

Fig. 20: Performance Analysis with GPU Configuration: Scratchpad=48K, L1-Cache=16K, Resident Threads=2048

scheduling policies respectively. We observe that *Shared-OWF-OPT* approach shows an average improvement of 17.73% and 18.08% with respect to unshared GTO and two-level scheduling policies respectively. The application *FDTD3d* degrade with our approach when compared to the baseline with either GTO scheduling or two-level scheduling since it has more number of L1 and L2 cache misses with sharing. The application *histogram* degrades with sharing when compared to the baseline with GTO scheduling because of more number of L1 misses. However *histogram* with sharing performs better than the baseline with two-level policy.

8.1.8. Comparison with Other GPGPU-Sim Configurations. We now compare the effectiveness of sharing approaches with *Unshared-LRR* for different GPGPU-Sim configurations. Figure 19 shows results for a GPU configuration that uses 48K L1 cache. We observe that sharing approach shows an average improvement of 14.04% with *Shared-OWF* and 18.71% with *Shared-OWF-OPT* over *Unshared-LRR*. The applications *DCT3*, *DCT4*, *SRAD1*, *SRAD2*, and *FDTD3d* using sharing are improved further with this configuration since they benefit from increased L1 cache size. For *heartwall*, *Unshared-LRR* benefits more with the increased L1 cache than the sharing approaches, it shows relatively less improvement of 86% when compared to Figure 14.

Figure 20 shows the performance comparison of sharing approaches with *Unshared-LRR* for a GPU configuration that has 48K scratchpad memory and the maximum number of resident threads in the SM as 2048. We observe average improvements of 8.62% and 9.21% with *Shared-OWF* and *Shared-OWF-OPT* approaches respectively. Consider the applications *DCT1*, *DCT2*, *DCT3*, and *DCT4*. With increase in the scratchpad memory, the number of resident thread blocks in the SM for these applications is not limited by the scratchpad memory, hence sharing does not increase the number of resident thread blocks. Also, the compiler optimizations do not insert the *relssp* instruction into their PTX code since there is no access to shared scratchpad region. Hence *Shared-OWF-OPT* behaves exactly same as *Shared-OPT*. The improvement in the performance of sharing approaches over *Unshared-LRR* is due to the OWF scheduling policy. OWF scheduler arranges the resident warps according to the owner warps. Since all the warps that are launched using *Shared-OWF* own their resources (no sharing), they become owner warps. Hence the warps are arranged according to their dynamic warp id, giving the observed benefit. While scratchpad sharing can increase the number of resident thread blocks for *SRAD1* and *SRAD2*, no additional blocks could be launched since the number of resident threads is restricted to 2048.

Figure 21 shows the performance comparison for a GPU configuration that uses 48K shared memory and the maximum number of resident threads as 3072. With the number of resident threads increasing from 2048 to 3072, sharing is able to increase the

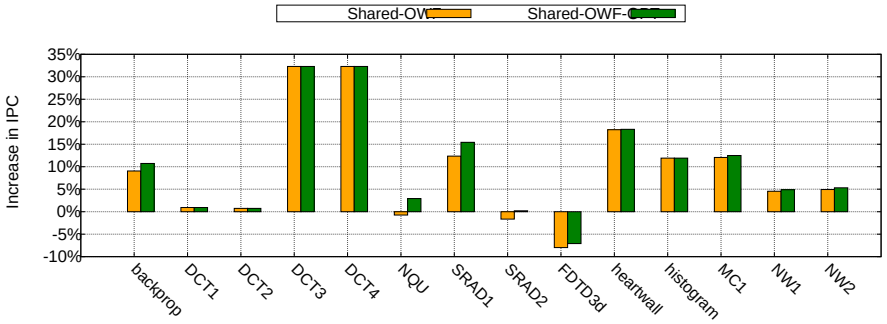

Fig. 21: Performance Analysis with GPU Configuration: Scratchpad=48K, L1-Cache=16K, Resident Threads=3072

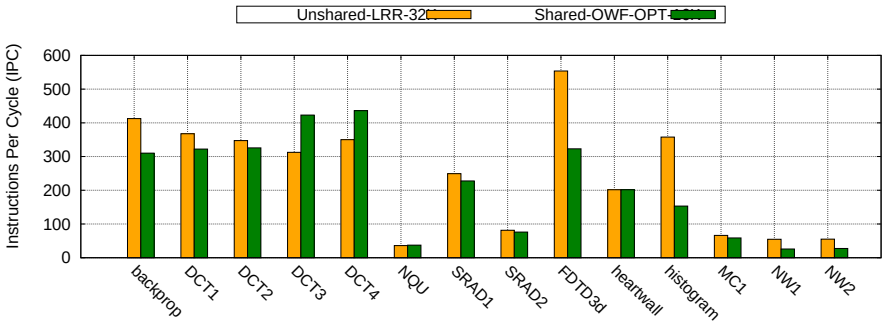

Fig. 22: Comparison with Unshared-LRR that Uses Twice the Scratchpad Memory

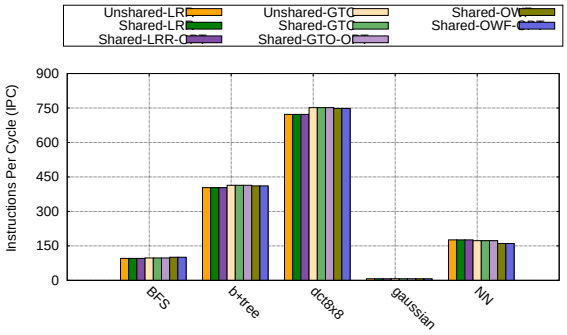

Fig. 23: Performance Comparison of Set-3 Benchmarks

number of resident thread blocks in *SRAD1* and *SRAD2* applications, thereby improving the performance.

8.1.9. Resource Savings. Figure 22 compares the IPC of *Shared-OWF-OPT* with *Unshared-LRR* that uses twice the amount of scratchpad memory on GPU. We observe that *DCT3*, *DCT4*, *NQU*, and *heartwall* show improvement with *Shared-OWF-OPT* over *Unshared-LRR* even with half the scratchpad memory. This is because sharing helps in increasing the TLP by launching additional thread blocks in each SM. The applications *DCT1*, *DCT2*, *SRAD1*, *SRAD2*, and *MC1* applications perform comparable with both the approaches. For the remaining applications, *Unshared-LRR* with double scratchpad memory performs better than sharing since more number of thread blocks are able to make progress with the former.

Table VII: Additional Benchmarks that are Limited by Scratchpad Memory

	Benchmark	Application	Kernel	#Scratchpad Variables	Scratchpad Size (Bytes)	Block Size
Benchmarks for 48KB / 64KB Scratchpad Memory						
1.	RODINIA	backprop	bpnn_layerforward_CUDA	2	9408	256
2.	CUDA-SDK	DCT1	CUDAkernel2DCT	1	8320	128
3.	CUDA-SDK	DCT2	CUDAkernel2IDCT	1	8320	128
4.	GPGPU-SIM	NQU	solve_nqueen_cuda_kernel	5	10496	64
5.	CUDA-SDK	histogram	histogram256Kernel	1	9216	192
6.	CUDA-SDK	marchingCubes (MC2)	generateTriangles	2	13824	48
7.	RODINIA	NW1	needle_cuda_shared_1	2	8452	32
8.	RODINIA	NW2	needle_cuda_shared_2	2	8452	32
Benchmarks for 48KB Scratchpad Memory						
9.	CUDA-SDK	FDTD3d	FiniteDifferencesKernel	1	3840	128
10.	RODINIA	heartwall	kernel	8	11872	128
11.	CUDA-SDK	marchingCubes (MC1)	generateTriangles	2	9216	32
Additional Benchmarks for 16KB Scratchpad Memory						
12.	RODINIA	kmeans	kmeansPoint	2	4608	576
13.	RODINIA	lud	lud_internal	2	3872	484

8.2. Analysis of Set-3 Benchmarks

Performance analysis of Set-3 benchmarks is shown in the Figure 23. Recall that the number of thread blocks launched by these applications is not limited by the scratchpad memory. We observe that the performance of the applications with *Unshared-LRR*, *Shared-LRR*, and *Shared-LRR-OPT* is exactly the same. For Set-3 applications all thread blocks are launched in unsharing mode. Hence *Shared-LRR* behaves exactly same as *Unshared-LRR*. Since these applications do not use any shared scratchpad memory, our compiler optimizations do not insert *relssp* instruction in their PTX code. Hence the number of instructions executed by the *Shared-LRR-OPT* approach is same as that of *Shared-LRR*. Similarly, we see that the performance of applications with *Unshared-GTO*, *Shared-GTO*, and *Shared-GTO-OPT* is exactly the same. However, with OWF optimization, *Shared-OWF* and *Shared-OWF-OPT* is comparable to the *Unshared-GTO* because OWF optimization arranges the resident warps according to the owner. Since all the the thread blocks own their scratchpad memory, they are sorted according to the dynamic warp id. Hence they perform comparable to *Unshared-GTO*. The performances with *Shared-OWF* and *Shared-OWF-OPT* are the same because the compiler optimizations do not insert any *relssp* instruction.

8.3. Additional Experiments

8.3.1. Performance Comparison with Other Configurations. To further verify the effectiveness of our approach, we evaluated it on two GPU configurations that use scratchpad memory of size 48KB and 64KB per SM (Table VIII). These configuration parameters are similar to that of NVIDIA's Kepler and Maxwell architectures respectively. The benchmarks that are used for the evaluation are shown in Table VII. The changes in Table VII with respect to those in Table I are:

— Kernel scratchpad memory size for *DCT1* and *DCT2* is increased from 2112 to 8320. This change ensures that applications are limited by scratchpad memory for both the configurations.
— A new application, *MC2*, is created based on *MC1*—the only difference being that kernel scratchpad memory size is increased to 13824, this is to enable scratchpad sharing for Configuration-2.
— The applications *DCT3*, *DCT4*, *SRAD1*, and *SRAD2* are dropped as scratchpad sharing could not be made applicable even by increasing the kernel scratchpad

Table VIII: GPGPU-Sim Configurations used for Additional Experiments

Resource/Core	Configuration-1	Configuration-2
Scratchpad Memory	48KB	64KB
Max Number of Thread Blocks	16	32
Max Number of Threads	2048	2048

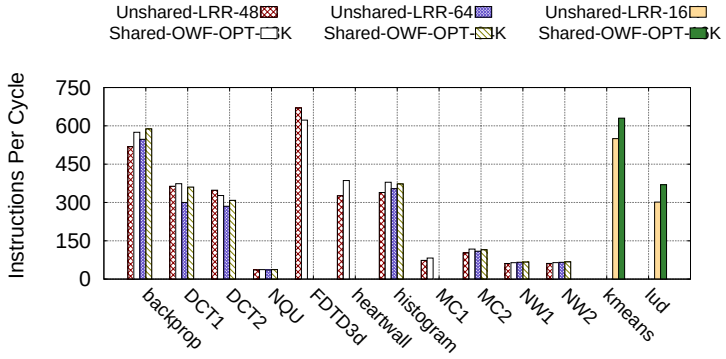

Fig. 24: Performance Analysis for Various Configurations

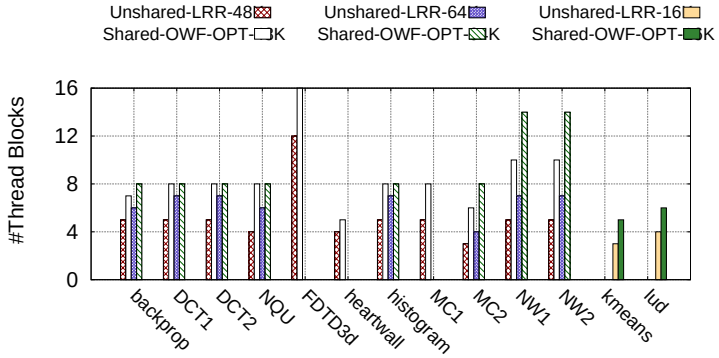

Fig. 25: No. of Thread Blocks for Various Configurations

memory size. These applications are limited either by the number of thread blocks or by the number of threads.

Note that scratchpad sharing is applicable for *MC1*, *FDTD3d*, and *heartwall* only with Configuration-1, but not with Configuration-2 (Table VIII). In addition, we performed experiments with 16KB scratchpad memory per SM for two new applications, *kmeans* and *lud*.

Increase in Number of Resident Thread Blocks: Figure 25 shows the increase in number of resident thread blocks for scratchpad sharing. In the figure, we use *Unshared-LRR-48K* to denote the baseline approach that uses 48KB scratchpad memory (according to Configuration-1) and *Shared-OWF-OPT-48K* to denote the scratchpad sharing approach (with all optimizations) that use 48KB scratchpad memory. The notations *Unshared-LRR-64K* and *Shared-OWF-OPT-64K*, which use 64KB scratchpad memory configuration, are defined analogously. From the figure, we observe that all the applications launch more number of resident thread blocks with our approaches (*Shared-OWF-OPT-48K* and *Shared-OWF-OPT-64K*) when compared to their respective baseline configurations. The applications *DCT1*, *DCT2*, and *FDTD3d* launch 8, 8, and 16

Table IX: Benchmarks used for comparison with Shared Memory Multiplexing [Yang et al. 2012]

Application	Kernel	#Scratchpad Variables	Scratchpad Size (Bytes)	Block Size
Convolution (CV)	convolutionColumnsKernel	1	8256	128
Fast Fourier Transform (FFT)	kfft	1	8704	64
Histogram (HG)	histogram256	1	7168	32
MarchingCubes (MC)	generateTriangles	2	9216	32
Matrix Vector Multiplication (MV)	mv_shared	1	4224	32
ScalarProd (SP)	scalarProdGPU	1	4114	64

Table X: Details of Modifications to the Shared Memory Multiplexing Benchmarks

Application	File Name	Line Number	Variable Name	Original Configuration	New Configuration
MC	marchingCubes.cpp	202	g_bQAReadback	flase	true
CV	main.cpp	73	iterations	10	1
SP	scalarProd_kernel.cu	42	accumResult (size)	ACCUM_N	ACCUM_N+12
HG	hist.cu	93	iterator	16	1
		43	-	runTest(1024, 1024)	runTest(1024, 896);
	hist_kernel.cu	3	BIN_SIZE	128	224
		4	THREAD	64	32
MV	mv.cu	107	iterator	16	1

thread blocks respectively with *Shared-OWF-OPT-48K*, which is the limit on the number of resident threads in the SM. Similarly *backprop*, *DCT1*, and *DCT2* launch 8 thread blocks each with *Shared-OWF-OPT-64K*, which is also the limit on the number of resident threads in the SM. For all the other applications, our approach is able to launch maximum number of thread blocks that is possible with scratchpad sharing approach. Hence, scratchpad sharing helps in increasing the number of thread blocks for the applications even with the configurations that use larger scratchpad memory per SM.

Figure 24 shows the performance comparison of our approach with the two baseline configurations. We observe that with *Shared-OWF-OPT-48K*, all the applications except *FDTD3d* and *DCT2* show performance improvement when compared to *Unshared-LRR-48*. Similarly with *Shared-OWF-OPT-64K*, all the applications perform better when compared to *Unshared-LRR-64K*. Consider the applications *backprop*, *DCT1*, *histogram*, *MC2*, and *NQU*. These applications with *Shared-OWF-OPT-48K* perform better even when compared to *Unshared-LRR-64K*. Also, these applications, when used with *Shared-OWF-OPT-64K*, perform better than *Unshared-LRR-64K* as well. *NW1* and *NW2* show improvement with scratchpad sharing when compared to their respective baseline approaches. Applications *heartwall* and *MC1*, where scratchpad sharing is applicable only with Configuration-1 (Table VIII), show improvement when compared to *Unshared-LRR-48K*. For *DCT2*, increasing the amount of scratchpad memory from 48KB to 64KB per SM does not improve the performance of *Unshared-LRR* configuration since it increases number of stall cycles in the SM. Hence increase in the number of thread blocks for *Shared-OWF-OPT-48K* does not improve its performance w.r.t *Unshared-LRR-48K*. However, *Shared-OWF-OPT-48K* performs better than *Shared-OWF-OPT-64K* because *Shared-OWF-OPT-48K* has lesser number of thread blocks that make more progress, hence reducing resource contention. *FDTD3d* does not show improvement with scratchpad sharing due to increase in the stall cycles with our approach. Also additional benchmarks, *kmeans* and *lud* show improvement with *Shared-OWF-OPT-16K* when compared to *Unshared-LRR-16K*.

To summarize, our approach helps in improving the performance of the applications even with increase in the size of the scratchpad memory per SM.

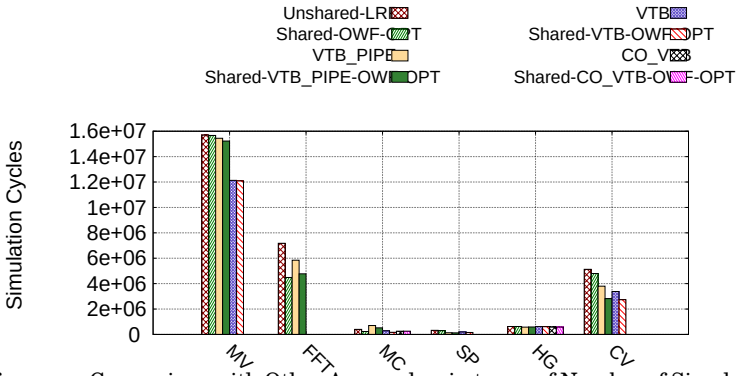

Fig. 26: Performance Comparison with Other Approaches in terms of Number of Simulation Cycles

Table XI: Comparing the Number of Simulated Instructions for Various Approaches

Benchmark	Unshared-LRR	Shared-OWF-OPT	VTB_PIPE	Shared-VTB_PIPE-OWF-OPT	VTB	Shared-VTB_OWF-OPT	CO_VTB	Shared-CO_VTB-OWF-OPT
MV	412,680,192	412,712,960	637,648,896	637,698,048	696,909,824	696,942,592	N/A	N/A
FFT	416,022,528	416,284,672	445,644,800	446,038,016	N/A	N/A	N/A	N/A
MC	24,321,552	24,387,088	29,228,560	29,294,096	27,352,592	27,418,128	25,456,176	25,521,712
SP	27,472,640	27,505,408	15,611,648	15,644,416	15,946,752	15,979,520	N/A	N/A
HG	19,898,368	19,906,560	10,682,368	10,690,560	10,430,464	10,438,656	13,253,518	13,261,710
CV	439,812,096	440,991,744	525,729,792	527,499,264	497,025,024	498,794,496	N/A	N/A

8.3.2. Performance Comparison with Shared Memory Multiplexing [Yang et al. 2012].

Figures 26 and 27 compare the performance of our approach with the software approaches proposed by Yang et al. [2012] in terms of number of simulation cycles and instructions per cycles respectively. We use their benchmarks (Table IX), and simulate them on the GPU configuration shown in Table II. In the figure, *Unshared-LRR* and *Shared-OWF-OPT* denote the baseline and scratchpad sharing approaches respectively. *VTB*, *VTB_PIPE*, and *CO_VTB* denote the compiler optimizations proposed by Yang et al. [2012][9]. Similarly, we use *Shared-VTB-OWF-OPT*, *Shared-VTB_PIPE-OWF-OPT*, and *Shared-CO_VTB-OWF-OPT* to measure performance of scratchpad sharing on the applications that are optimized with *VTB*, *VTB_PIPE*, and *CO_VTB* respectively.

From Figure 26 we observe that the application *MC* performs better (spends less number of simulation cycles) with *Shared-OWF-OPT* than with *Unshared-LLR*, *VTB*, *VTB_PIPE*, and *CO_VTB* approaches. Interestingly, applying *Shared-OWF-OPT* on top of *VTB*, *VTB_PIPE*, or *CO_VTB* improves the performance further. Similarly, *FFT* shows improvement with *Shared-OWF-OPT* when compared to *Unshared-LRR* and *VTB_PIPE*. In this case also *Shared-VTB_PIPE-OWF-OPT* outperforms VTB_PIPE. In contrast, for the application *HG*, sharing does not impact the performance, even on the top of *VTB* and *VTB_PIPE* optimizations. This is because the additional thread blocks launched do not make much progress before they start accessing shared scratchpad. For the same reason, scratchpad sharing does not have any further impact on *MV*. The applications *CV* and *SP* perform better with *VTB_PIPE* than with *Shared-OWF-OPT*. However, the performance is further improved when scratchpad sharing is combined with *VTB* and *VTB_PIPE* approaches. We also observe a change in the number of executed instructions for *VTB*, *VTB_PIPE*, and *CO_VTB* approaches for all the applica-

[9] Note that, as described in Yang et al. [2012], *CO_VTB* is suitable only for few workloads (i.e., *MC* and *HG*). Also, for *FFT*, we do not compare *VTB* with *VTB_PIPE* because *VTB* combines 4 thread blocks whereas *VTB_PIPE* combines 2 thread blocks in their implementation.

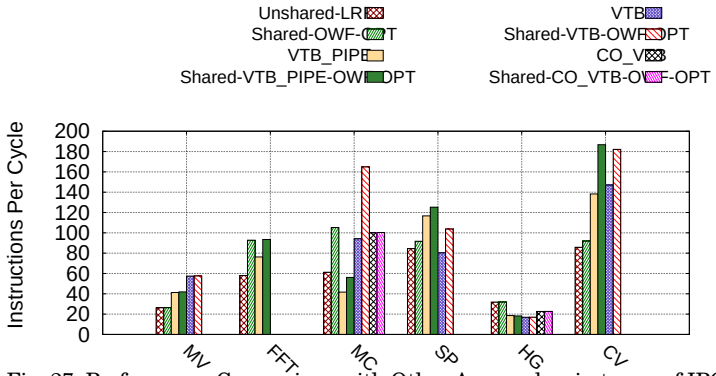

Fig. 27: Performance Comparison with Other Approaches in terms of IPC

tions (shown in Table XI) when compared the *Unshared-LRR*, because they modify the benchmarks.

It can be concluded from these experiments that scratchpad sharing and shared memory multiplexing approaches compliment each other well, and most applications show the best performance when the two approaches are combined.

8.3.3. Performance Analysis for Different Number of SMs. Figure 28 shows the effectiveness of our approach when different number of SMs are used in GPUs. Table XII shows five different configurations of SMs (obtained from GPGPUSIM [2014]) that specify the number of clusters and number of SMs in each cluster. In the figure *Unshared-LRR-1* and *Shared-OWF-OPT-1* represent the baseline and scratchpad sharing approaches when they are simulated with *SM Config-1*. Similarly, we use other notations to represent the approaches for other configurations.

From the figure, we observe that applications show maximum benefit with baseline approach when the number of SMs is 30. However, the applications show further improvements when scratchpad sharing is applied on them. Application *heartwall* shows maximum benefit with scratchpad sharing, because the additional thread blocks launched in our approach do not access shared scratchpad memory. For *backprop*, the performance improves with increase in the number of SMs, since increase in the number of SM leads to increase in TLP. The performance improves further when scratchpad sharing is applied. *MC1* also shows improvement with scratchpad sharing for all configurations when compared to baseline, because the additional thread blocks launched in our approach make significant progress before they start accessing shared scratchpad memory. *DCT1, DCT2, DCT3, and DCT4* perform better when the number of SMs is 14 (*SM Config-1*) when compared to 15 (*SM Config-2*). This is because *SM Config-2* contains more number of SMs per clusters than *SM Config-1*, and all the SMs within a cluster share a common port with interconnect. *NQU* shows similar improvements for all configurations with our approach. Because even with increase in the number of SMs, baseline approach does not show improvement.

From the results we conclude that, scratchpad sharing helps in improving the performance applications even by varying number of SMs in a GPU. The applications improve further when there are more number of SMs with less number of SMs within a cluster in GPUs.

9. RELATED WORK

Resource sharing technique [Jatala et al. 2016], proposed by the authors earlier, improves the throughput by minimizing the register and scratchpad memory underuti-

Table XII: GPGPU-Sim Configurations with Various Number of SMs

Resource/Core	SM Config-1	SM Config-2	SM Config-3	SM Config-4	SM Config-5
Total Number of SMs	14	15	16	16	30
Number of Cluster * Number of SMs per Cluster	7*2	3*5	8*2	4*4	10*3

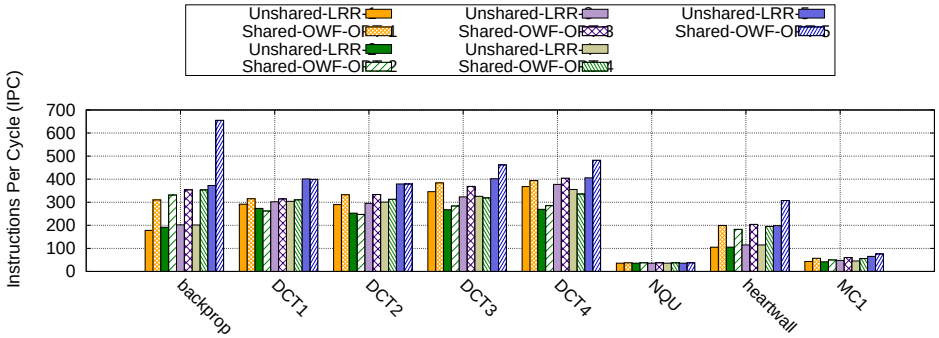

Fig. 28: Performance Analysis for Different Number of SMs

lization by modifying the GPU architecture and scheduling algorithm. This work improves on it by introducing compiler optimizations for better layout of scratchpad variables and early release of shared scratchpad. Other related approaches that improve the performance of GPUs are discussed below:

9.1. Resource Management in GPUs

Shared memory multiplexing [Yang et al. 2012] proposes solutions that come closest to our approach. They provide software and hardware solutions to address the TLP problem caused by limited shared memory. The software approach combines two thread blocks into a single virtual thread block. The two thread blocks in a virtual block can execute instructions in parallel, as long as they do not access shared memory; and become serial when they need to access shared memory. The paper also describes a mechanism (called CO-VTB) that divides the shared memory into private and public part so that the thread blocks in a virtual block can access the private part in parallel and the public part in serial. However, CO-VTB has a high overhead of partitioning the data into private and public part, and is not suitable for all workloads. Also, they need to generate the code manually. The paper also gives a hardware solution to dynamically allocate and deallocate scratchpad memory using the existing barrier instruction. Again, these instructions need to be inserted manually in the code, and nesting of the barrier instructions is not allowed in order to avoid any deadlocks.

In contrast, ours is a hardware solution that allows launching additional thread blocks in each SM. These additional thread blocks use the wasted scratchpad memory, and also share part of the allocated scratchpad memory with other resident thread blocks. The additional thread blocks launched in our approach make progress as long as they do not require shared scratchpad memory, and wait until the shared scratchpad is released by other thread blocks. Our approach is fully automatic—our compiler optimization automatically identifies the regions of the shared and unshared scratchpad memory at run-time and inserts instruction to release the shared scratchpad as early as possible. Even in the presence of barrier instructions, our approach can not have deadlocks. In addition, we propose a warp scheduling mechanism that effectively schedule these additional warps to hide the long latencies in a better way.

Warp level divergence technique [Xiang et al. 2014] improves the TLP by minimizing register underutilization. It launches one additional partial thread block when there are insufficient number of registers for an entire thread block. However, the number of warps in the partial thread block is decided by the number of unutilized registers, and also the partial thread block does not share registers with any other thread blocks. The unified storage approach [Gebhart et al. 2012] allocates the resources of SM (such as registers, scratchpad memory, and cache) dynamically as per the application demand. Tarjan and Skadron [2011] use virtual registers to launch more thread blocks. These registers are mapped to the physical registers as per the demand. Our compiler optimizations can help in early release of unused registers with this approach. Gomez-Luna et al. [2013] describe a mechanism to lock and unlock parts of scratchpad memory. We can reutilize the existing mechanism by defining a custom hash function that maps shared scratchpad memory regions to corresponding lock addresses. For unshared scratch region, the access can be given directly (i.e., the lock can be granted always). Kayiran et al. [2013] propose a dynamic algorithm to launch the optimal number of thread blocks in an SM to reduce the resource contention. We can combine their techniques with our approach to reduce the number of increase in the stall cycles can occur with the additional thread blocks. Li et al. [2011] propose a resource virtualization scheme for sharing of GPU resources with multiprocessors. The virtualization layer proposed by them helps in improving the performance by overlapping multiple kernels executions.

Our approach is different from the above in that it addresses underutilization of scratchpad memory by launching more than one additional thread blocks in each SM. These thread blocks share scratchpad memory with other resident thread blocks, thus improve its utilization. The proposed OWF-scheduler reduces stall cycles and improves the performance further.

9.2. Compiler Optimizations for Efficient Resource Utilization in GPUs

Ma and Agrawal [2010] formulated the problem of scratchpad memory allocation as an integer programming problem, which maximizes scratchpad memory access and minimizes device memory access to improve GPU performance. Their framework can allocate parts of arrays on scratchpad, and also suggest profitable loop transformations. Hayes and Zhang [2014] proposed on-chip memory allocation scheme for efficient utilization of GPU resources. It aims to alleviate register pressure by spilling registers to scratchpad memory instead of local memory. Xie et al. [2015] proposed a compile time coordinated register allocation scheme to minimize the cost of spilling registers. These schemes do not propose any architectural change to GPUs and are orthogonal to our approach of scratchpad sharing.

9.3. Scheduling Techniques for GPUs

The two level warp scheduling algorithm, proposed by Narasiman et al. [2011], forms groups of warps and uses LRR to schedule warps in a group. It also proposes a large warp microarchitecture to minimize resource underutilization. Lee and Wu [2014] hide the long execution latencies by scheduling critical warps more frequently than other than warps. It helps in finishing the thread block sooner thus improving resource utilization. However, it requires the knowledge of critical warps. To address the problem, Lee et al. [2015] proposed a coordinated solution that identifies the critical warps at run time using instructions and stall cycles. Further, they proposed a greedy based critical warp scheduling algorithm to accelerate the critical warps in the SMs. OWL [Jog et al. 2013] provides a scheduling mechanism to reduce cache contention and to improve DRAM bank level parallelism. Lee et al. [2014] focus on reducing resource contention by providing lazy thread block scheduling mechanism. They also

proposed block level CTA scheduling policy that allocates consecutive CTAs into the same SM to exploit cache locality. Their approach can also be integrated to our approach.

9.4. Improving GPU Performance through Memory Management

Several other approaches exploit memory hierarchy to improve the performance of GPU applications. Li et al. [2015b] proposed compiler techniques to efficiently place data onto registers, scratchpad memory, and global memory by analyzing data access patterns. Sethia et al. [2015] proposed a scheduling policy that improves the GPU performance by prioritizing memory requests of single warp when memory saturation occurs. Li et al. [2015a] provide a mechanism to handle the cache contention problem that occurs due to increased number of resident threads in an SM. Their approach is alternative to the earlier proposed thread throttling techniques [Rogers et al. 2013, 2012; Kayiran et al. 2013].

9.5. Problems with Warp Divergence

Other techniques to improve GPU performance is by handling warp divergence. Dynamic warp formation [Fung et al. 2007] addresses the limited thread level parallelism that is present due to branch divergence. It dynamically forms new warps based on branch target condition. However, the performance of this approach is limited by the warp scheduling policy. Thread block compaction [Fung and Aamodt 2011] addresses the limitation of dynamic warp formation that occurs when the new warps that are formed may require more number of memory accesses. Their approach provides a solution by regrouping the new warps at the reconverging points. However in their solution, warps need to wait for other warps to reach the divergent path. Anantpur and Govindarajan [2014] proposed linearization technique to avoid duplicate execution of instructions that occurs due to branch divergence in GPUs. Brunie et al. [2012]; Han and Abdelrahman [2011] provide hardware and software solutions to handle branch divergence in GPUs.

9.6. Miscellaneous

Warped pre-execution [Lee et al. 2016] accelerates a single warp by executing independent instructions when a warp is stalled due to long latency instruction. It improves the GPU performance by hiding the long latency cycles in a better way. Baskaran et al. [2008] proposed a compiler framework for optimizing memory access in affine loops. Huo et al. [2010]; Gutierrez et al. [2008] show that several applications are improved by using scratchpad memory instead of using global memory.

10. CONCLUSIONS AND FUTURE WORK

In this paper, we propose architectural changes and compiler optimizations for sharing scratchpad effectively to address the underutilization of scratchpad memory in GPUs. Experiments with various benchmarks help us conclude that if the number of resident thread blocks launched by an application are limited by scratchpad availability (Table I), scratchpad sharing (with the compiler optimizations) improves the performance. On the other hand, for other applications where the number of thread blocks is not limited by scratchpad availability (Table IV), the hardware changes do not negatively impact the run-time.

In future, we would like to extend our work to integrate register sharing approach [Jatala et al. 2016]. Value range analysis techniques [Harrison 1977; Quintao Pereira et al. 2013], typically employed for detecting buffer overflows, can be incorporated in our approach to refine the access ranges of shared scratchpad vari-

ables, thus help release shared scratchpad even earlier. We need to study the impact of hardware changes on power consumption, and find ways to minimize it.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their useful feedback. We also thank the authors of "Shared Memory Multiplexing: A Novel Way to Improve GPGPU Throughput" [Yang et al. 2012] for sharing their benchmarks with us.

REFERENCES

Jayvant Anantpur and R. Govindarajan. 2014. Taming Control Divergence in GPUs through Control Flow Linearization. In *Compiler Construction (CC)*.

A Bakhoda, G.L. Yuan, W.W.L. Fung, H. Wong, and T.M. Aamodt. 2009. Analyzing CUDA workloads using a detailed GPU simulator. In *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*.

Muthu Manikandan Baskaran, Uday Bondhugula, Sriram Krishnamoorthy, J. Ramanujam, Atanas Rountev, and P. Sadayappan. 2008. A Compiler Framework for Optimization of Affine Loop Nests for GPGPUs. In *International Conference on Supercomputing (ICS)*.

Nicolas Brunie, Sylvain Collange, and Gregory Diamos. 2012. Simultaneous Branch and Warp Interweaving for Sustained GPU Performance. In *International Symposium on Computer Architecture (ISCA)*.

Shuai Che, M. Boyer, Jiayuan Meng, D. Tarjan, J.W. Sheaffer, Sang-Ha Lee, and K. Skadron. 2009. Rodinia: A benchmark suite for heterogeneous computing. In *IEEE International Symposium on Workload Characterization (IISWC)*.

CUDA 2012. CUDA C Programming Guide. (2012). http://docs.nvidia.com/cuda/pdf/CUDA_C_Programming_Guide.pdf

CUDA-SDK 2014. CUDA-SDK. (2014). http://docs.nvidia.com/cuda/cuda-samples

Gregory Frederick Diamos, Andrew Robert Kerr, Sudhakar Yalamanchili, and Nathan Clark. 2010. Ocelot: a dynamic optimization framework for bulk-synchronous applications in heterogeneous systems. In *Parallel Architectures and Compilation Techniques (PACT)*.

Wilson W. L. Fung and Tor M. Aamodt. 2011. Thread Block Compaction for Efficient SIMT Control Flow. In *High Performance Computer Architecture (HPCA)*.

Wilson W. L. Fung, Ivan Sham, George Yuan, and Tor M. Aamodt. 2007. Dynamic Warp Formation and Scheduling for Efficient GPU Control Flow. In *International Symposium on Microarchitecture (MICRO)*.

Mark Gebhart, Stephen W. Keckler, Brucek Khailany, Ronny Krashinsky, and William J. Dally. 2012. Unifying Primary Cache, Scratch, and Register File Memories in a Throughput Processor. In *International Symposium on Microarchitecture (MICRO)*.

Juan Gomez-Luna, Jose Maria Gonzalez-Linares, Jose Ignacio Benavides Benitez, and Nicolas Guil. 2013. Performance Modeling of Atomic Additions on GPU Scratchpad Memory. *IEEE Trans. Parallel Distrib. Syst.* (2013).

GPGPUSIM 2014. GPGPU-Sim Simulator. (2014). http://www.gpgpu-sim.org

Eladio Gutierrez, Sergio Romero, MariaA. Trenas, and EmilioL. Zapata. 2008. Memory Locality Exploitation Strategies for FFT on the CUDA Architecture. In *International Meeting on High-Performance Computing for Computational Science (VECPAR)*.

Tianyi David Han and Tarek S. Abdelrahman. 2011. Reducing Branch Divergence in GPU Programs. In *Workshop on General Purpose Processing on Graphics Processing Units (GPGPU)*.

William H. Harrison. 1977. Compiler Analysis of the Value Ranges for Variables. *IEEE Trans. Software Eng.* 3, 3 (1977).

Ari B. Hayes and Eddy Z. Zhang. 2014. Unified On-chip Memory Allocation for SIMT Architecture. In *International Conference on Supercomputing (ICS)*.

Xin Huo, V.T. Ravi, Wenjing Ma, and G. Agrawal. 2010. Approaches for parallelizing reductions on modern GPUs. In *High Performance Computing (HiPC)*.

Vishwesh Jatala, Jayvant Anantpur, and Amey Karkare. 2015. The More We Share, The More We Have: Improving GPU performance through Register Sharing. *CoRR* abs/1503.05694 (2015).

Vishwesh Jatala, Jayvant Anantpur, and Amey Karkare. 2016. Improving GPU Performance Through Resource Sharing. In *High-Performance Parallel and Distributed Computing (HPDC)*.

Adwait Jog, Onur Kayiran, Nachiappan Chidambaram Nachiappan, Asit K. Mishra, Mahmut T. Kandemir, Onur Mutlu, Ravishankar Iyer, and Chita R. Das. 2013. OWL: Cooperative Thread Array Aware Scheduling Techniques for Improving GPGPU Performance. In *Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.

John B. Kam and Jeffrey D. Ullman. 1976. Global Data Flow Analysis and Iterative Algorithms. *J. ACM* 23, 1 (Jan. 1976).

O. Kayiran, A Jog, M.T. Kandemir, and C.R. Das. 2013. Neither more nor less: Optimizing thread-level parallelism for GPGPUs. In *Parallel Architectures and Compilation Techniques (PACT)*.

Uday Khedker, Amitabha Sanyal, and Bageshri Karkare. 2009. *Data Flow Analysis: Theory and Practice* (1st ed.). CRC Press, Inc., Boca Raton, FL, USA.

Minseok Lee, Seokwoo Song, Joosik Moon, J. Kim, Woong Seo, Yeongon Cho, and Soojung Ryu. 2014. Improving GPGPU resource utilization through alternative thread block scheduling. In *High Performance Computer Architecture (HPCA)*.

Sangpil Lee, Won Woo Ro, Keunsoo Kim, Gunjae Koo, Myung Kuk Yoon, and Murali Annavaram. 2016. Warped-preexecution: A GPU pre-execution approach for improving latency hiding. In *High Performance Computer Architecture (HPCA)*.

Shin-Ying Lee, Akhil Arunkumar, and Carole-Jean Wu. 2015. CAWA: Coordinated Warp Scheduling and Cache Prioritization for Critical Warp Acceleration of GPGPU Workloads. In *International Symposium on Computer Architecture (ISCA)*.

Shin-Ying Lee and Carole-Jean Wu. 2014. CAWS: Criticality-aware Warp Scheduling for GPGPU Workloads. In *Parallel Architectures and Compilation Techniques (PACT)*.

Chao Li, Yi Yang, Zhen Lin, and Huiyang Zhou. 2015b. Automatic data placement into GPU on-chip memory resources. In *Code Generation and Optimization (CGO)*.

Dong Li, Minsoo Rhu, Daniel R. Johnson, Mike O'Connor, Mattan Erez, Doug Burger, Donald S. Fussell, and Stephen W. Redder. 2015a. Priority-based cache allocation in throughput processors. In *High Performance Computer Architecture (HPCA)*.

T. Li, V. K. Narayana, E. El-Araby, and T. El-Ghazawi. 2011. GPU Resource Sharing and Virtualization on High Performance Computing Systems. In *International Conference on Parallel Processing (ICPP)*.

Wenjing Ma and Gagan Agrawal. 2010. An Integer Programming Framework for Optimizing Shared Memory Use on GPUs. In *Parallel Architectures and Compilation Techniques (PACT)*.

Steven S. Muchnick. 1997. *Advanced Compiler Design and Implementation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

Veynu Narasiman, Michael Shebanow, Chang Joo Lee, Rustam Miftakhutdinov, Onur Mutlu, and Yale N. Patt. 2011. Improving GPU Performance via Large Warps and Two-level Warp Scheduling. In *International Symposium on Microarchitecture (MI-*

CRO).

Sreepathi Pai, R. Govindarajan, and Matthew J. Thazhuthaveetil. 2014. Preemptive Thread Block Scheduling with Online Structural Runtime Prediction for Concurrent GPGPU Kernels. In *Parallel Architectures and Compilation Techniques (PACT)*.

PTX 2014. Parallel Thread Execution. (2014). http://docs.nvidia.com/cuda/parallel-thread-execution/

Fernando Magno Quintao Pereira, Raphael Ernani Rodrigues, and Victor Hugo Sperle Campos. 2013. A Fast and Low-overhead Technique to Secure Programs Against Integer Overflows. In *Code Generation and Optimization (CGO)*.

Timothy G. Rogers, Mike O'Connor, and Tor M. Aamodt. 2012. Cache-Conscious Wavefront Scheduling. In *International Symposium on Microarchitecture (MICRO)*.

Timothy G. Rogers, Mike O'Connor, and Tor M. Aamodt. 2013. Divergence-aware Warp Scheduling. In *International Symposium on Microarchitecture (MICRO)*.

Ankit Sethia, D Anoushe Jamshidi, and Scott Mahlke. 2015. Mascar: Speeding up GPU warps by reducing memory pitstops. In *High Performance Computer Architecture (HPCA)*.

D. Tarjan and K. Skadron. 2011. On demand register allocation and deallocation for a multithreaded processor. http://www.google.com/patents/US20110161616. (2011). US Patent App. 12/649,238.

Ping Xiang, Yi Yang, and Huiyang Zhou. 2014. Warp-level divergence in GPUs: Characterization, impact, and mitigation. In *High Performance Computer Architecture (HPCA)*.

Xiaolong Xie, Yun Liang, Xiuhong Li, Yudong Wu, Guangyu Sun, Tao Wang, and Dongrui Fan. 2015. Enabling Coordinated Register Allocation and Thread-level Parallelism Optimization for GPUs. In *International Symposium on Microarchitecture (MICRO)*.

Yi Yang, Ping Xiang, Mike Mantor, Norm Rubin, and Huiyang Zhou. 2012. Shared Memory Multiplexing: A Novel Way to Improve GPGPU Throughput. In *Parallel Architectures and Compilation Techniques (PACT)*.

A. APPENDIX

Table XIII: Comparing IPC values of Scratchpad sharing with Various Baseline Implementations

Benchmark	Unshared-LRR	Unshared-GTO	Unshared-2Level	Shared-OWF-OPT
backprop	178.01	179.78	178.38	310.1
DCT1	284.48	289.39	281.54	322.28
DCT2	283.84	287.16	288.42	325.83
DCT3	358.11	373.21	384.3	423.12
DCT4	381.23	402.4	406.12	436.2
NQU	35.77	35.77	35.76	37.46
SRAD1	199.18	200.58	199.23	227.74
SRAD2	67.19	67.39	66.98	76.18
FDTD3d	330.52	331.96	328.95	322.94
heartwall	104.92	104.92	105.05	201.62
histogram	153.46	153.38	151.56	153.19
MC1	44.43	44.57	44.15	58.79
NW1	25.34	25.34	25.26	25.94
NW2	25.4	25.4	25.32	27.51

In Figure 14, we have shown the performance of scratchpad sharing approach when normalized with respect to *Unshared-LRR*. In Table XIII, we show the absolute number of instructions executed per cycle (IPC) for scratchpad sharing approach, and we compare it with that of baseline implementation that uses LRR, GTO, and two-level scheduling policies.